Methods of Computational Neuroscience

Part 5 of 6

Marine Biological Laboratory - Woods Hole, MA
August 2009

Day 5 (August 14)

Nancy Kopell

43:01

53:09

Sara A. Solla

123:21

- MBL09_SOLLA.ppt
- PRL04.pdf
- Chaos07.pdf

Week 3 (Aug. 17-21)

Day 1 (Aug. 17)

Karel Svoboda

- Svoboda_MCN.pdf

Larry Abbott

- AbbottLect1.pdf
- force5.pdf
- Bi, G.-q. & Poo, M.M. Synaptic modification by correlated activity: Hebb's postulate revisited. Annu. Rev. Neurosci. 24,139-166.
- Caporale N. & Dan, Y. Spike Timing-Dependent Plasticity: A Hebbian LearningRule. Annu. Rev. Neurosci. 31, 25-46 (2008).
- Dan Y. & Poo M.M. Spike timing-dependent plasticity: from synapse to perception. Physiol Rev. 86, 1033-1048 (2006).

- mageeNatRevNeuro00.pdf

Bartlett Mel

- Mel_MCN_lecture.pdf

Day 4 (August 13)

Emre Aksay

- MCNTalk_2009_Aksay_Wiki.pptx

Mark Goldman

- Goldman_MCNTalk_Theory09.pdf

Day 2 (August 18)

Sara Solla

Haim Sompolinsky:

- Learning spike based rules-THE TEMPOTRON
- guetig_sompolinsky06_full.pdf
- gutig_sompolinsky_plos09.pdf

Larry Abbott's second talk: Chaotic networks

- AbbottTalk2.pdf
    - Jaeger H., and Haas H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science 304, 78-80.

Day 3 (August 19)

Bence Olveczky

- woods hole09.pdf

Ila Fiete

- 2009_aug_MBL_birdsong_topost.pdf

Day 4 (August 20)

Reza Shadmehr

- Woods_hole_09_2.ppt

Emo Todorov

- MBL09.pdf

Day 5 (August 21)

Simon Laughlin

- mcn 09 simon.pdf

Week 4 (Aug. 24-28)

Day 1 (Aug. 24)

Mitya Chklovskii

- Theoretical Connectomics.pdf

Problem Sets

Week 1

Simulated Neurons

Problem Set 1.zip
See the papers from day 1 (Fairhall et al., 2006 and Pillow et al., 2008)

Information theory

ProblemSet-InformationTheory.pdf

ProblemSet-FisherInfoCorrelatedNeurons-MCN2009.pdf (updated 2009.08.11)
Solutions
ProblemSetSolution-InformationTheory-MCN2009.pdf (updated 2009.08.11)
ProblemSetSolution-FisherInfoCorrelatedNeurons-MCN2009.pdf
(updated 2009.08.11)

Spectral analysis

Link to Chronux download website

Kleinfeld_Mitra_CSHL_2009_revised..pdf
matlab_for mcn_2990.m
example_data.mat
whisking.mat
090606_PA5_noStim.mat

Week 2

Working with XPP

NOTE: These problems are to be done at your discretion, and will not be presented.

- tcurrent.ode Exploring T-current and bifurcations
- ping.ode A simple model for PING

These are text files. The instructions are at each individual file after the 'Done' statement.

NEURON tutorial and exercises

- NEURON Tutorial ppt.pdf
- NEURON short exercises.pdf
  - Complete these exercises to become familiar with NEURON, as well as gaining a better understanding of dendritic dynamics.
- simple_layerVpyramidal.zip

- Webpage for downloading NEURON to your computer
  - Follow the instructions for your OS of choice. At the computer stations in the lab we use the MacOS-X 10.4 version
  - For Mac laptops, you HAVE to install the XCode package (size is about 1GB). Please ask the TA's for it, before the tutorial on Wednesday.

NEURON mini projects

- NEURON mini projects.pdf
- Rall_MCN_Assignment.pdf
- Katz et al 2009 - Neuron.pdf
- Kampa & Stuart 2006 - JNeurosci.pdf

This version of the code contains an additional .hoc file for use with the second mini project. Open "exerciseS2.hoc" to access it.

mbl0813.zip

for problem 3:
Calcium spikes in basal dendrites (Kampa and Stuart 2006) [modelDB]

Attractor network exercises

Mark Goldman will discuss your solutions with you on Monday afternoon.

MCN_RecurrentNetsExercise.pdf

MCN 2009 Lecture Notes & Papers
Reader – Part 1 of 6

Week 1 (Aug. 3-7)

Day 1 (Aug. 3) - Reader – Part 1 of 6

Adrienne Fairhall

- mcn09FairhallSlides.pdf
- Fairhall2006.pdf (Covariance analysis in retina)
- Pillow2008.pdf (GLM analysis in retina)
- Fairhall-etal-01.pdf (Multiple timescale adaptation in the fly)
- lundstromNN08.pdf (Fractional differentiation in cortical neurons)
- MaravallPLOS07.pdf (Gain control in barrel cortex)
- ruderman1997.pdf (Review of natural scene statistics)

Bill Bialek (website)

- bialek_02.pdf Les Houches: *Thinking about the brain*
- The ideas about optimal motion estimation go back to a paper with Marc Potters, which can be found here: http://www.princeton.edu/~wbialek/our_papers/potters+bialek_94.pdf
- The material on photon counting in vision is covered in a set of lecture notes which can be found at

  http://www.princeton.edu/~wbialek/PHY562.html

MCN 2009 Lecture Notes & Papers
Reader – Part 2 of 6

Day 2 (Aug. 4) - Reader – Part 2 of 6

Jonathan Pillow

- pillow_slides09_MaxLi_GLM.pdf (lecture slides)

Michael Berry (website)

- MCN 2009 Populations.ppt
- Recording All Cells.pdf
- Redundancy in Retinal Code.pdf

Rava da Silveira

- TheOrchestralBrain-WithAnnotations.pdf

Day 3 (Aug. 5)

Roozbeh Kiani

- Kiani_CategorizationInIT.pdf
- kianiJNP07.pdf
- kianiNeuron08.pdf

Jim DiCarlo

- DiCarloWHslides.pdf

Elad Schneidman

- Schneidman-MBL-MCN-2009.pdf
- Schneidman_2006_Nature_Weak pairwise correlations imply strongly.pdf
- Schneidman_PRL_2003_Network information and connected correlations.pdf
- Tkacik+al_06.pdf

Day 4 (Aug. 6)

Ehud Ahissar

- Ahissar_MCN2009_partA.pdf
- Ahissar_MCN2009_partB.pdf
- Knutsen & Ahissar TINS 2009 + Erratum.pdf
- Diamond et al NRN 2008.pdf

David Kleinfeld (website)

- MCN_2009_1.pdf
- MCN_2009_2_online.pdf
- john_curtis_NN_touch_S1 .pdf
- david_Encyclopedia_of_Neuroscience.pdf
- david_ehud_mat_Current_Opinions_06.pdf

Day 5 (Aug 7)

Bill Spain

- MCN 09short.ppt
- Reading list for Bill Spain's seminar.doc
- Jeffress.pdf (one of the papers on the reading list that is not available online)

Week 2 (Aug. 10-14)

Day 1 (Aug 10)

Bard Ermentrout

- chapt1.pdf
- A very comprehensive tutorial on bifurcations for neuro
- A dated tutrial for XPP
- A very quick look at using xpp to solve the HH eqns
- A link to the xpp homepage

ml.ode Morris-Lecar ODE

hhh.ode Hodgkin-Huxley

Installation of XPP on your computer:

- Using XPP on windows

Adrienne Fairhall

- mcn2009Fairhall_2.pdf
- hongAgueraFairhallNC07.pdf
- lundstromFamulareFairhall.pdf
- agueraFairhallBialek.pdf
- hongLundstromFairhallPloSCB08.pdf
- lundstromNC08.pdf

Day 2 (Aug 11)

Dan Johnston

Day 3

Michael Hausser

- Hausser_slides.pdf
- Sjostrom_Hausser_Neuron06.pdf
- Hausser&London2005review.pdf

Patterns of Neural Activity in Networks with Complex Connectivity

Sara Solla, Alex Roxin,
Santiago Madruga, Hermann Riecke

NORTHWESTERN
UNIVERSITY

Roxin, Riecke, Solla - Phys. Rev.Lett, **92**, 198101 (2004)
Riecke, Roxin, Madruga, Solla - Chaos **17**, 026110 (2007)

Form and Function

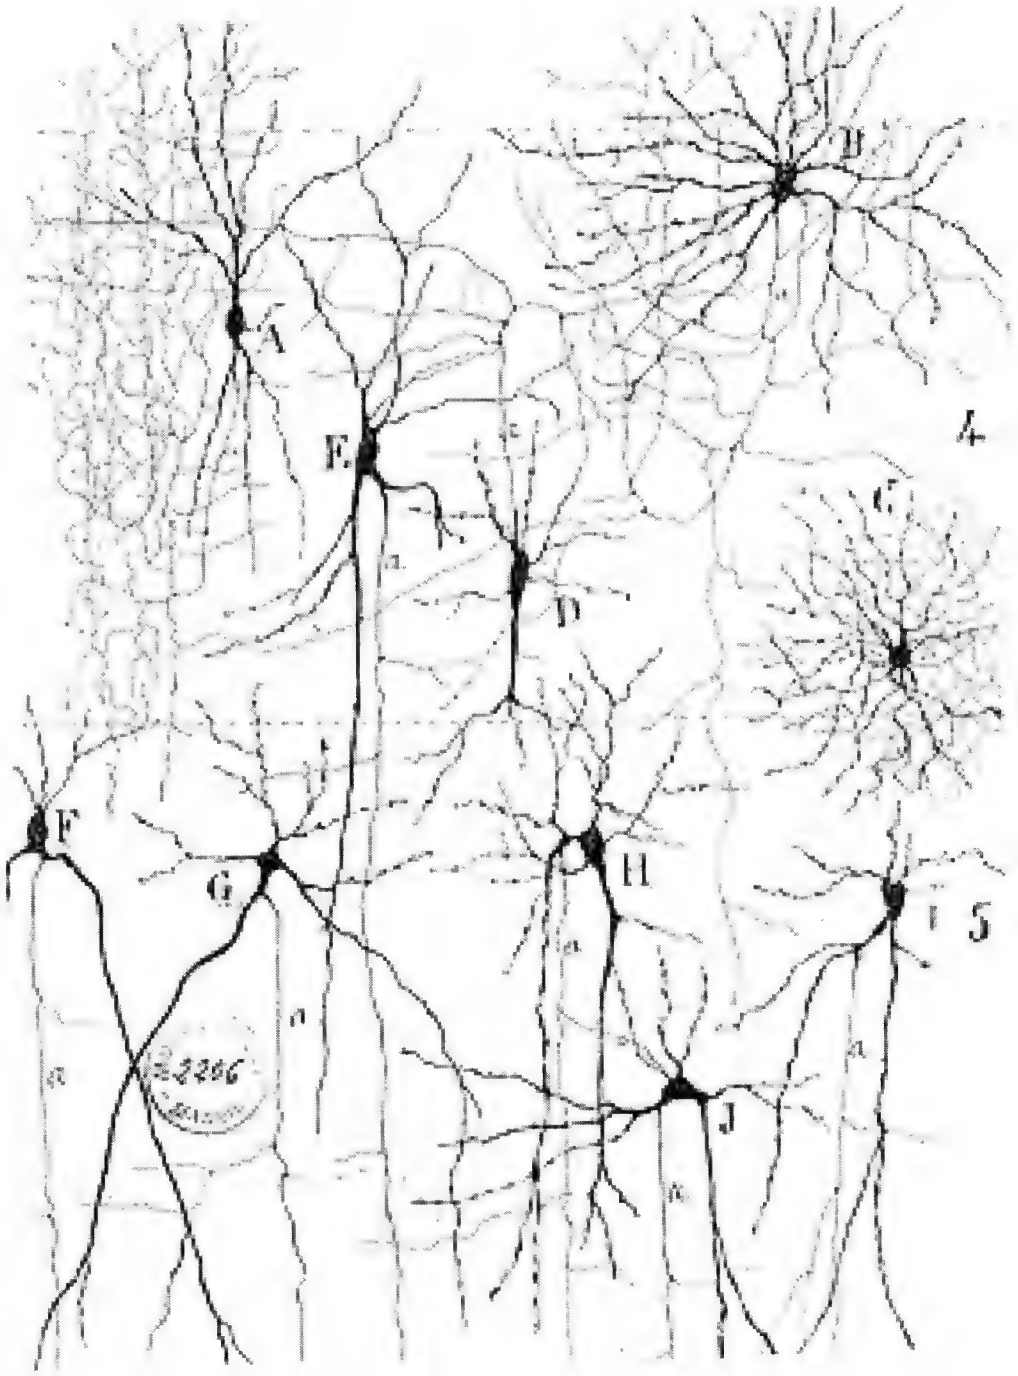

Form and Function

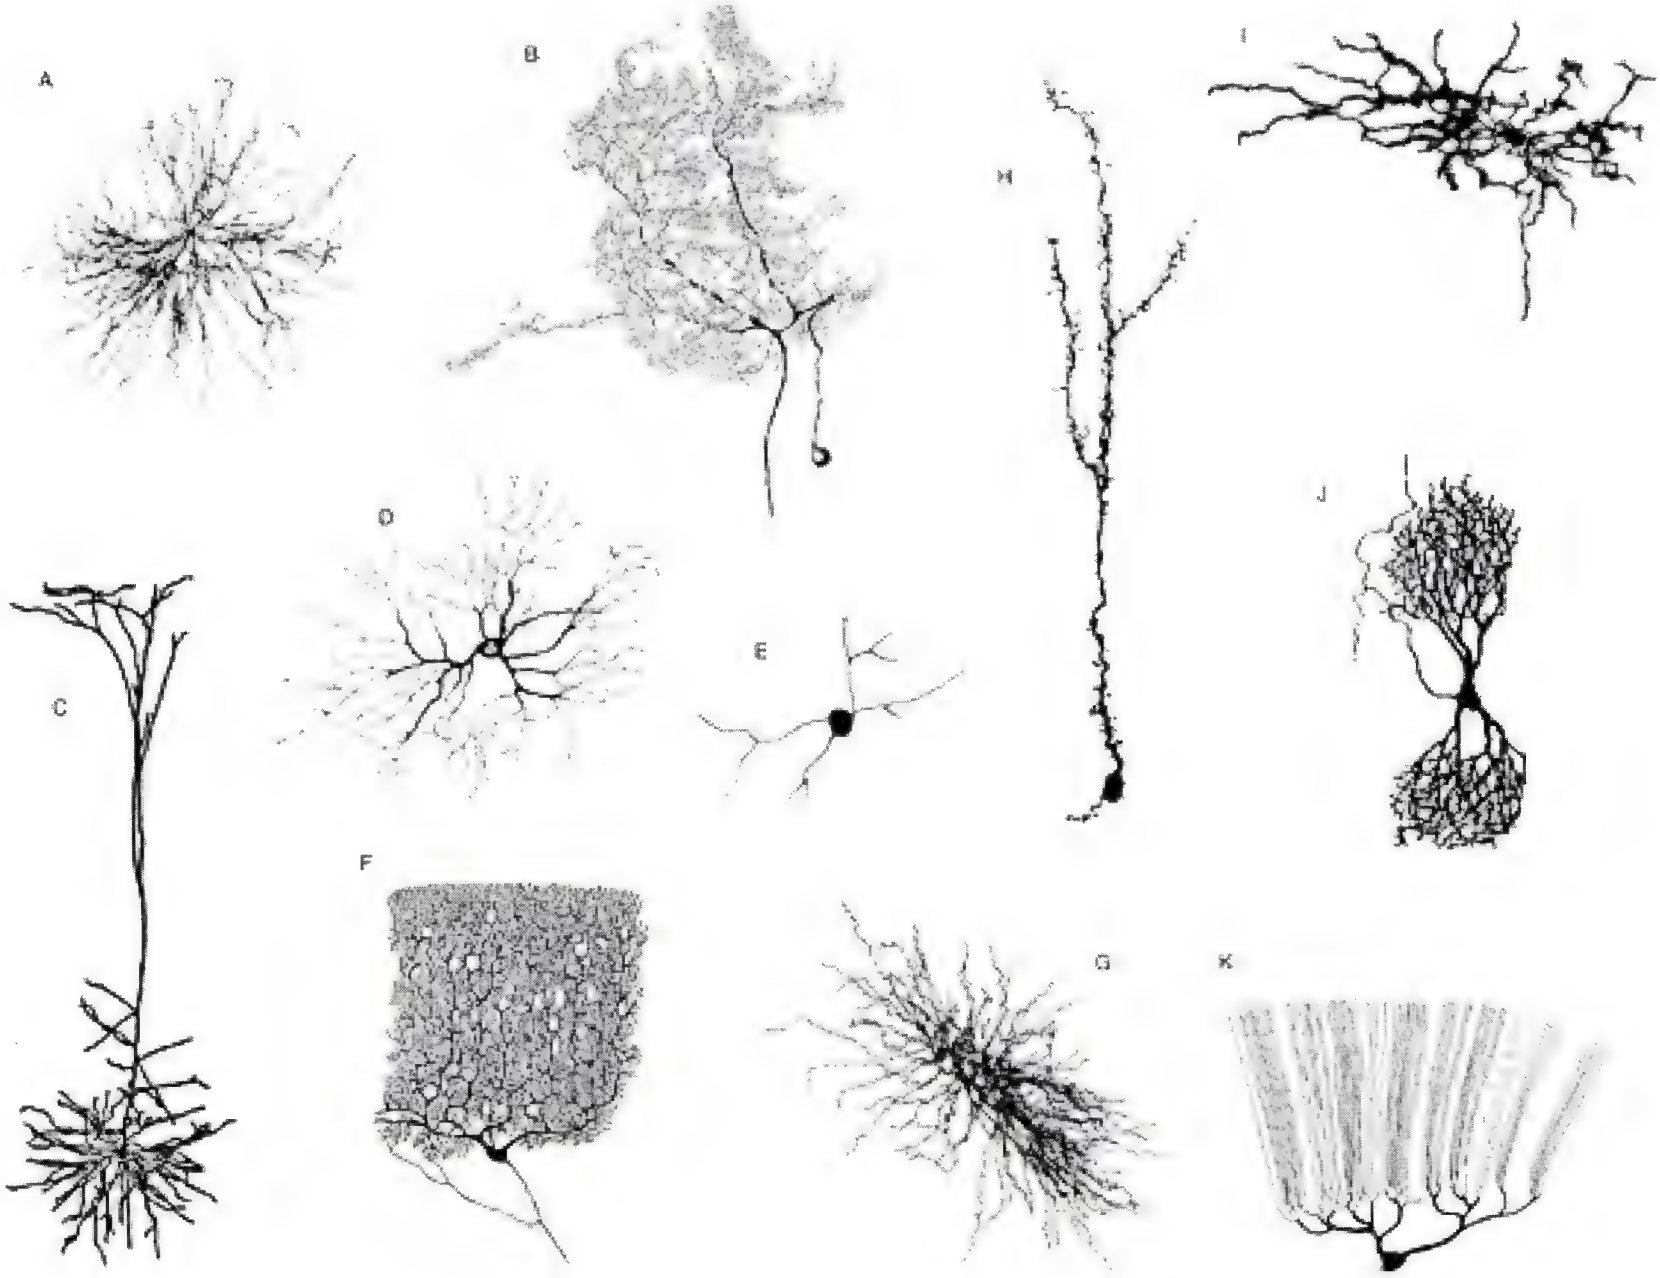

Form and Function

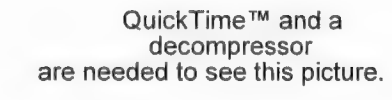

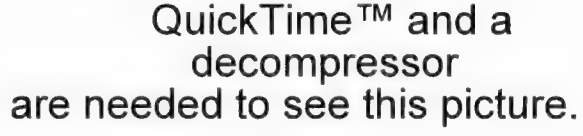

Network: Connectivity Diagram

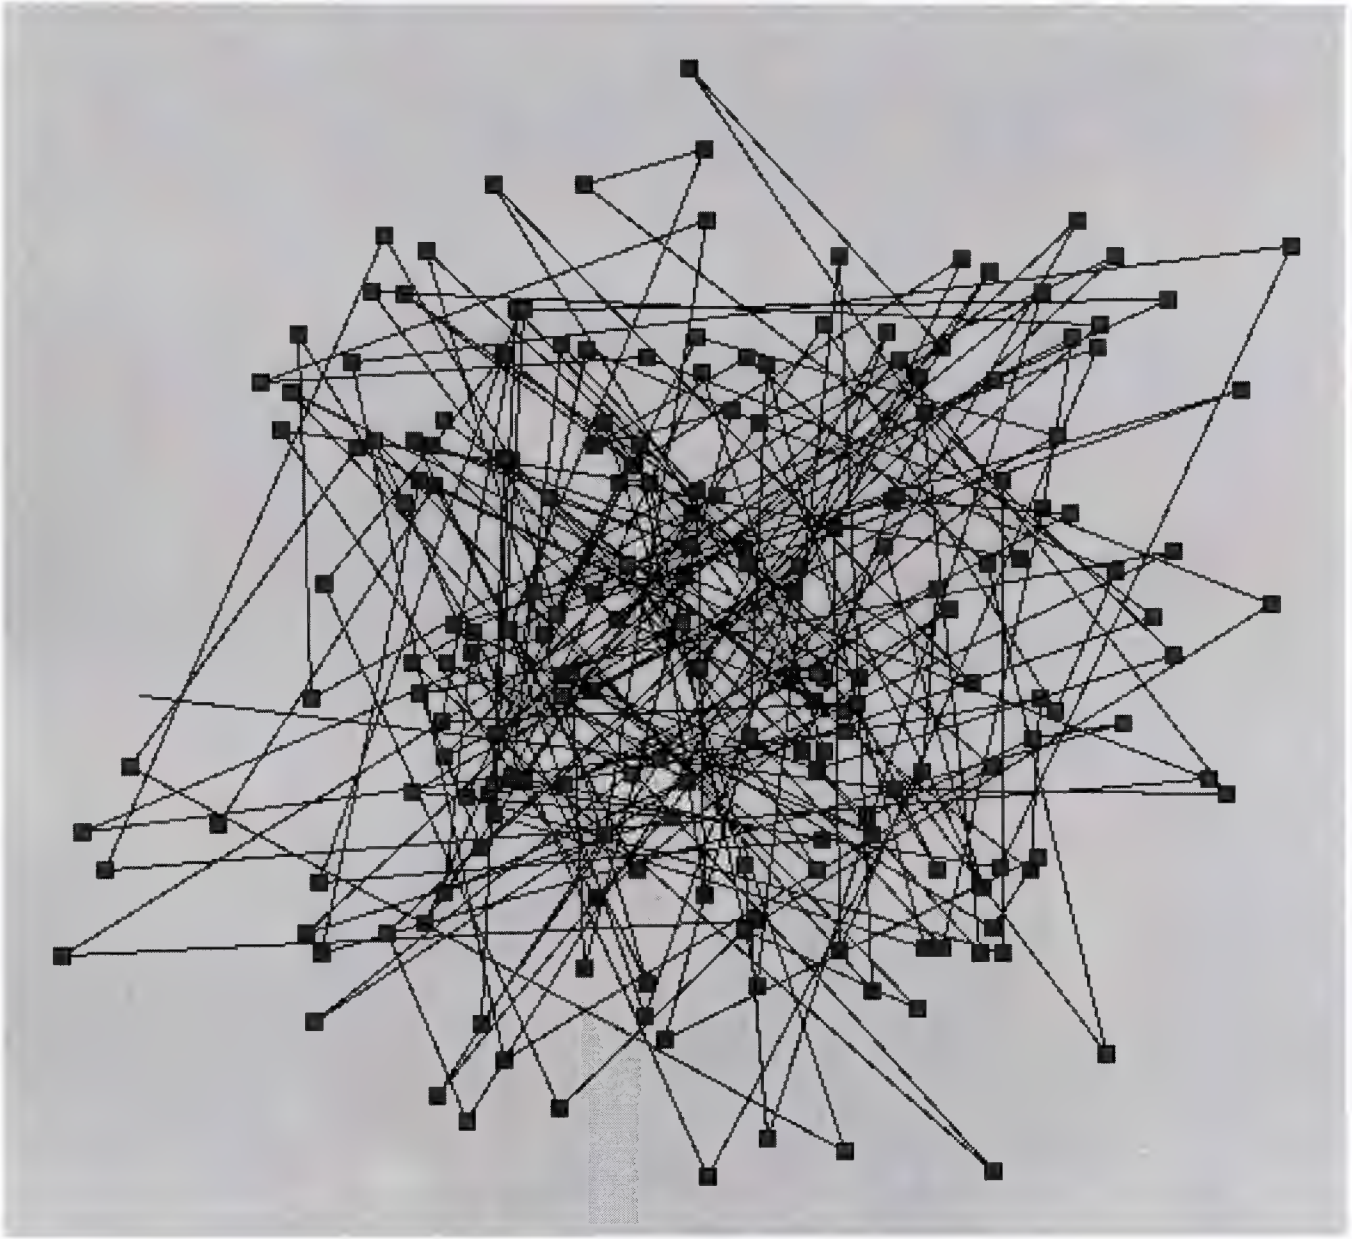

To which extent does network topology determine or affect network function?

Networks: Regular vs Random

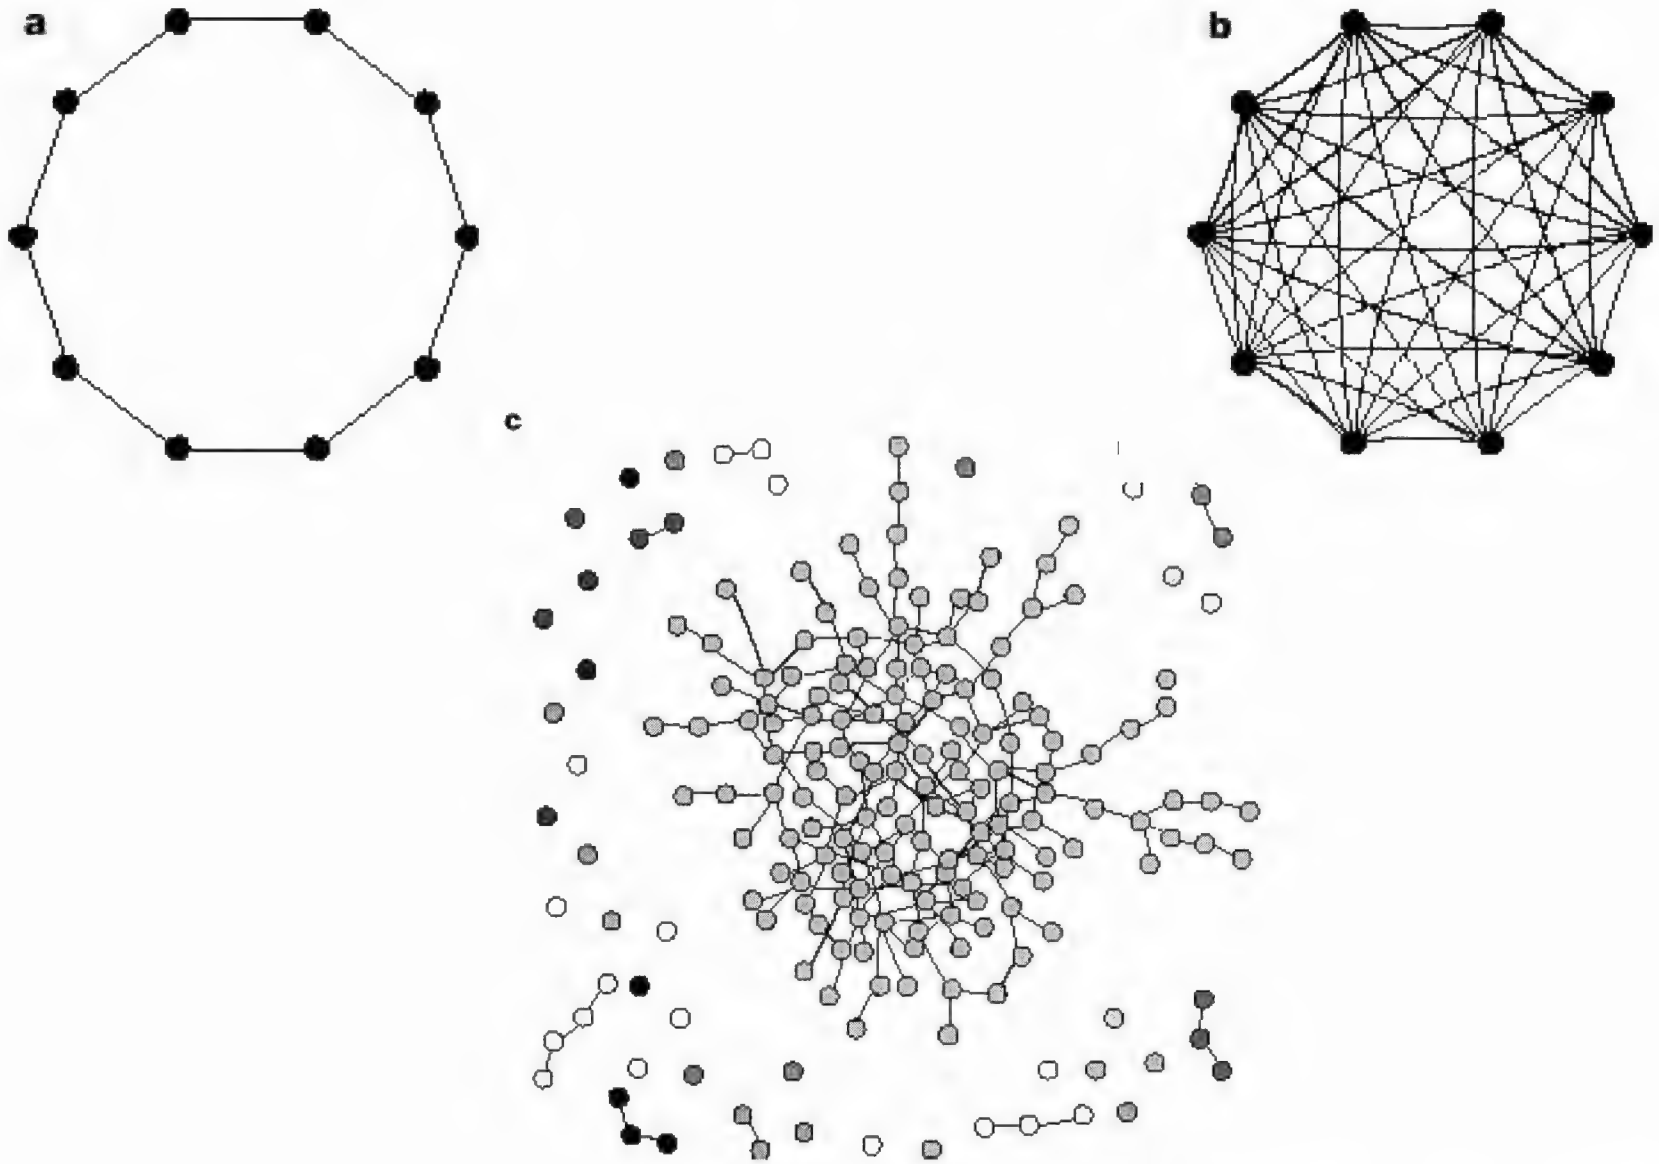

Strogatz - Nature **410,** 268 (2001

Networks: Complex

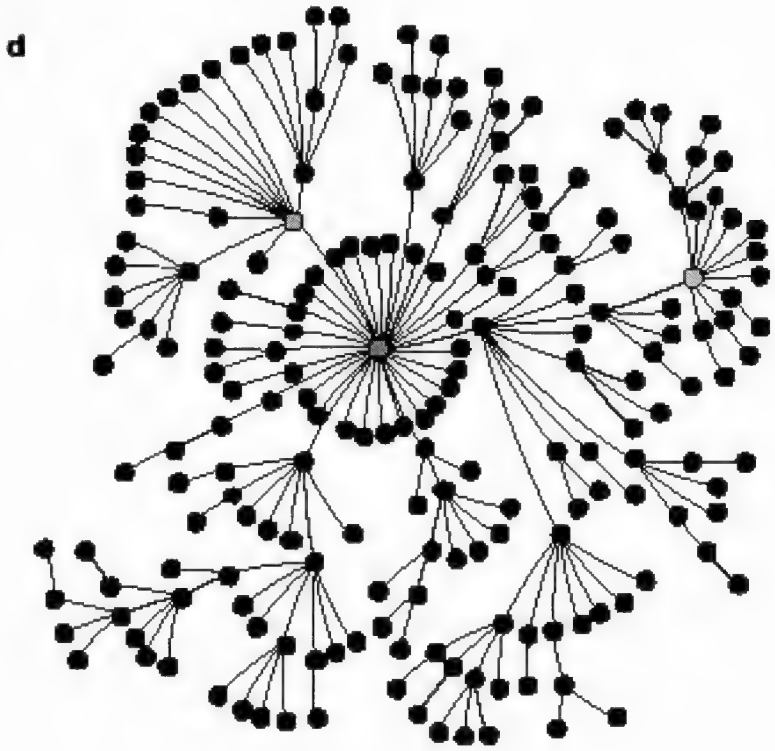

Scale Free

QuickTime™ and a
decompressor
are needed to see this picture.

Small World

Model of Neural Connectivity: a Small-World Network

Many complex networks have a small-world topology characterized by dense local clustering or cliquishness of connections between neighboring nodes yet a short path length between any (distant) pair of nodes due to the existence of relatively few long-range connections.

This is an attractive model for the organization of brain anatomical and functional networks because a small-world topology can support both segregated (specialized) and distributed (integrated) information processing.

Bassett, Bullmore - The Neuroscientist **12**, 512

Small-World Cortical Networks

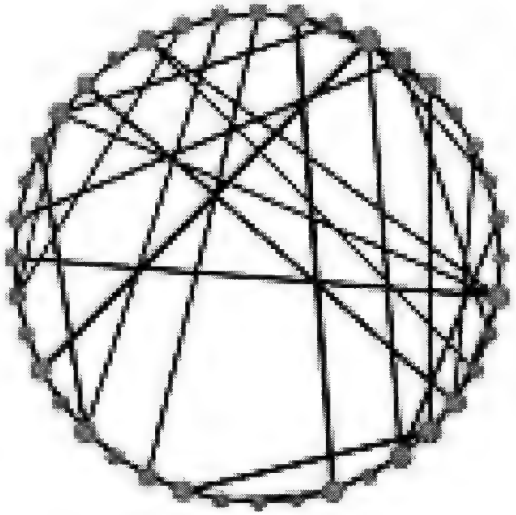

Glial cells provide a matrix of regular topology; local connectivity through calcium waves.

Neurons embedded in this structure can create intermediate and long-range links.

Neurons and glial cells can communicate with each other via calcium waves.

The aggregate system of neurons and glial cells can be viewed as a small-world network of excitable cells. Neurons, outnumbered by glial cells by one order of magnitude, provide the shortcuts.

Sinha, Saramaki, Kaski - Phys. Rev. E **76**, 015101

Small-World Brain Networks

•Activity in hippocampal slices has been successfully modeled using SW networks of excitatory neurons that reproduce both bursts (CA3) and seizures (CA1) [Netoff, Clewley, Arno, Keck, White - J. Neurosci. **24**, 8075 (2004)].

•Large-scale synchronization associated with epileptic seizures has been modeled using SW networks of Hindmarsh-Rose neurons [Percha, Szakpasu, Zochowski, Parent - Phys. Rev. E **72**, 031909 (2005)].

•A large-scale structural model of the dentate gyrus has been formulated and used to identify topological determinants of epileptogenesis [Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**, 1566 (2007)].

• Small-world networks are increasingly being applied to the analysis of human brain functional networks derived from electroencephalography, magnetoencephalography, and fMRI experiments [Eguiluz, Chialvo, Cecchi, Baliki, Apkarian - Phys. Rev. Lett. **94**, 018012 (2005)].

Dentate Gyrus: a Small-World Network

Complex network topology: neither regular, nor random

L : average path length C : clustering coefficient

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

The average number of synapses between any two neurons in the dentate gyrus is less than three - similar to the average path length for the nervous system of C. Elegans, which has only 302 neurons as opposed to over one million!

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**,

Model?

A SMALL-WORL NETWORK OF INTEGRATE-AND-FIRE NEURONS

Integrate-and-Fire Neurons

$$C_m \frac{dV}{dt} = I - g_L (V - E_L)$$

$$\tau \frac{dV}{dt} = -(V - V_\infty)$$

$$\tau = C_m / g_L \qquad V_\infty = E_L + I / g_L$$

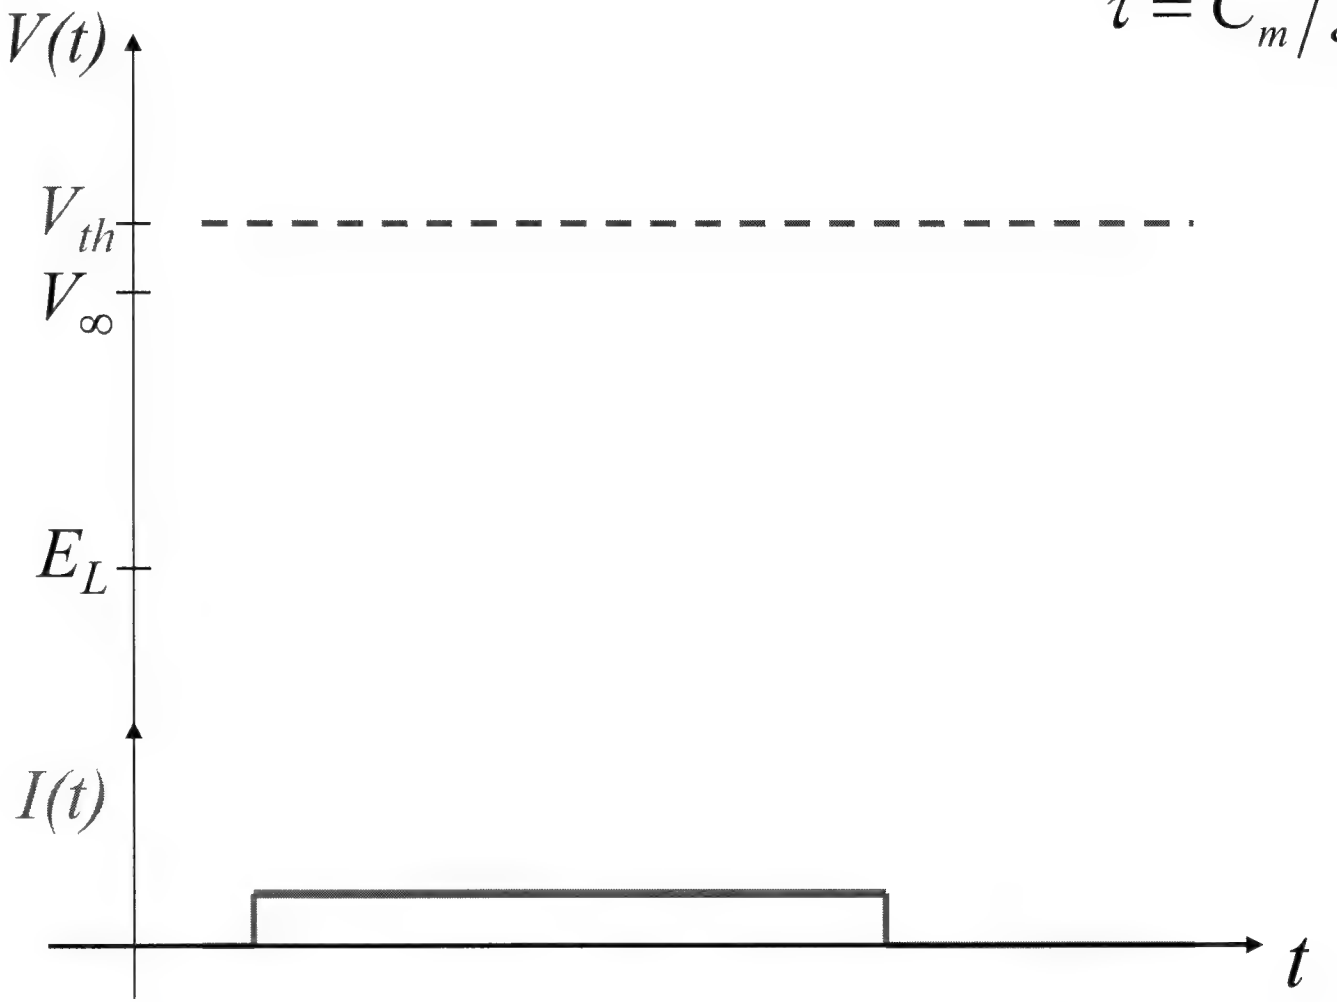

Excitable Neuron

$$V_\infty < V_{th}$$

Integrate-and-Fire Neurons

$$C_m \frac{dV}{dt} = I - g_L (V - E_L)$$

$$\tau \frac{dV}{dt} = -(V - V_\infty)$$

$$\tau = C_m / g_L \qquad V_\infty = E_L + I / g_L$$

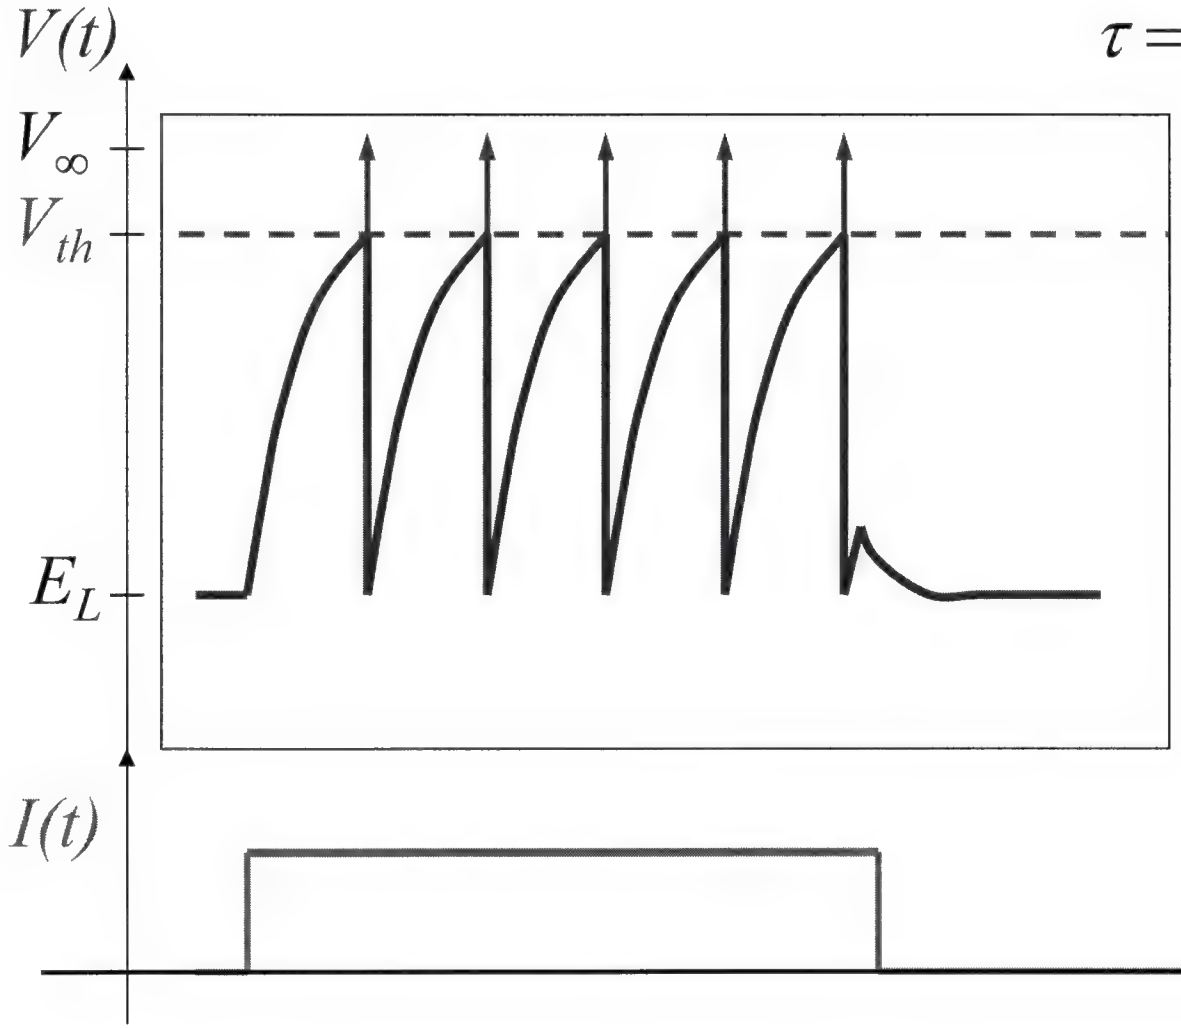

Periodically Firing Neuron

$$V_\infty > V_{th}$$

A Network of Excitable Integrate-and-Fire Neurons

$$\tau \frac{dV_i}{dt} = -V_i + V_\infty + V_{syn} \sum_{j=1}^{N} w_{ij}\, \delta(t - t^j - \tau_D)$$

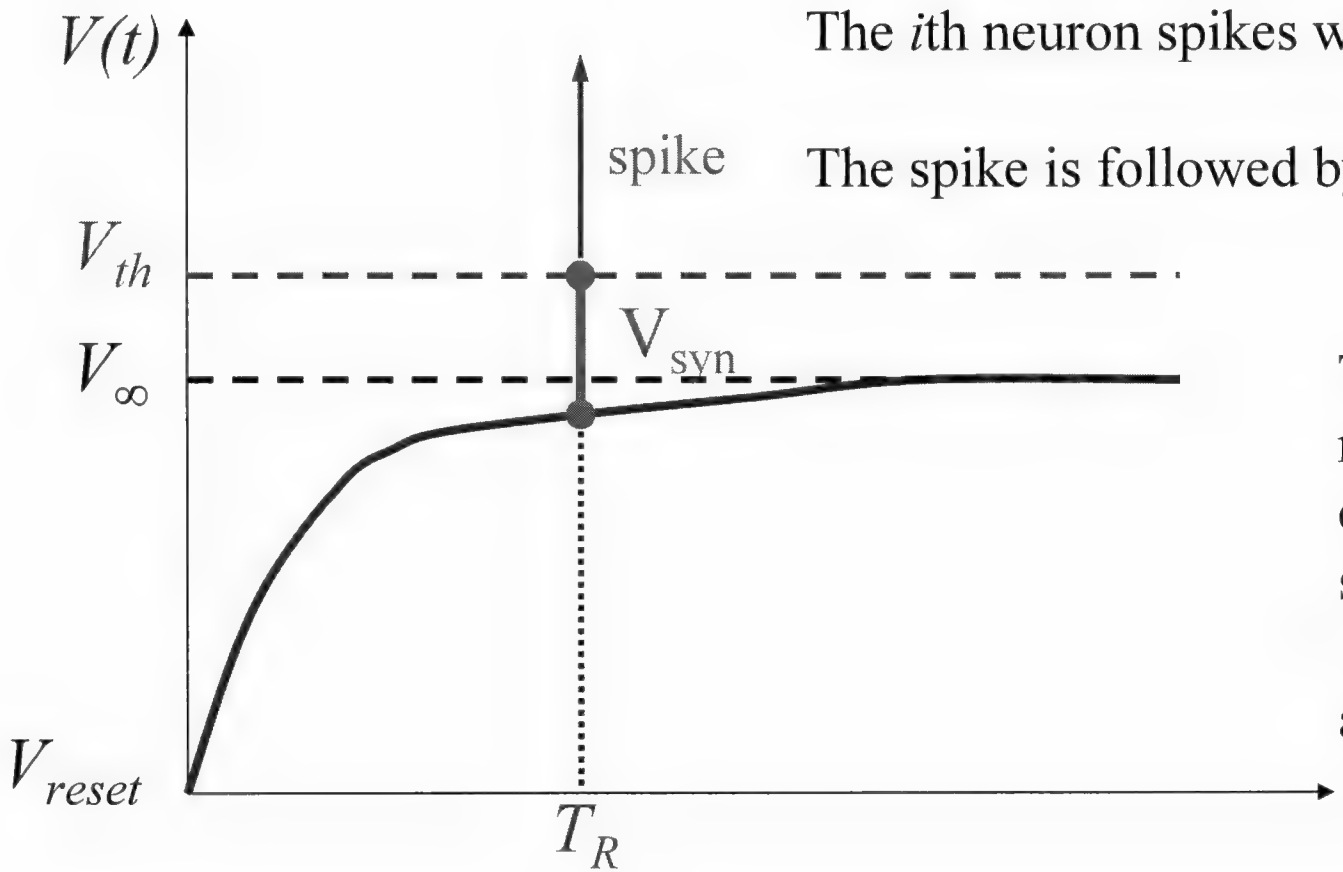

The ith neuron spikes whenever $V_i(t^-) = V_{th}$

The spike is followed by a reset $V_i(t^+) = V_{reset} = E_L$

The recovery time T_R is the minimum time that must elapse after a reset before a synaptic input of strength V_{syn} can make the cell fire again.

Sustained Network Activity: $p=0$

Sustained Network Activity: $p=0.05$

Sustained Network Activity

p=0

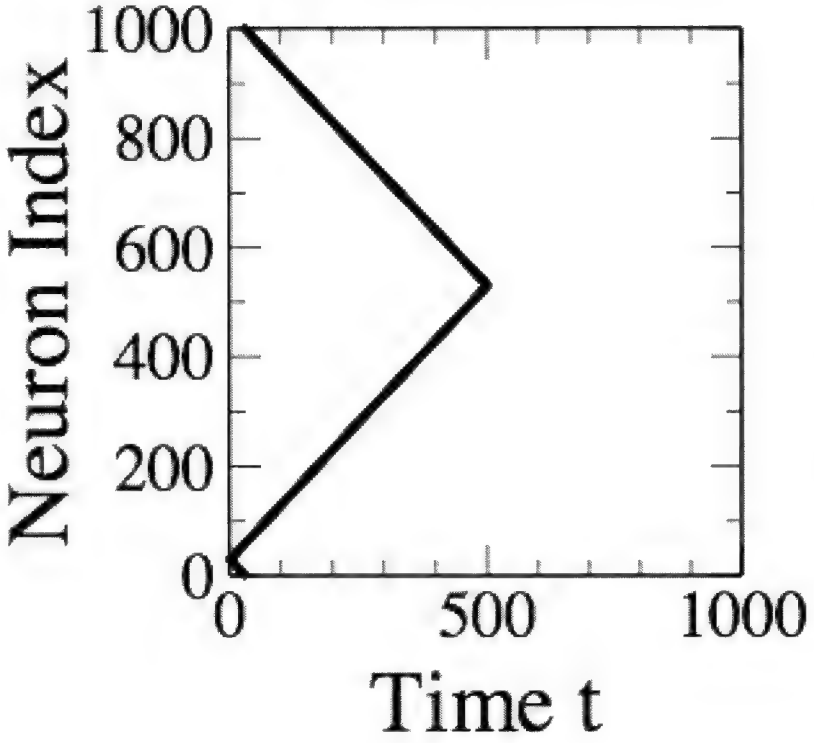

p=0.05

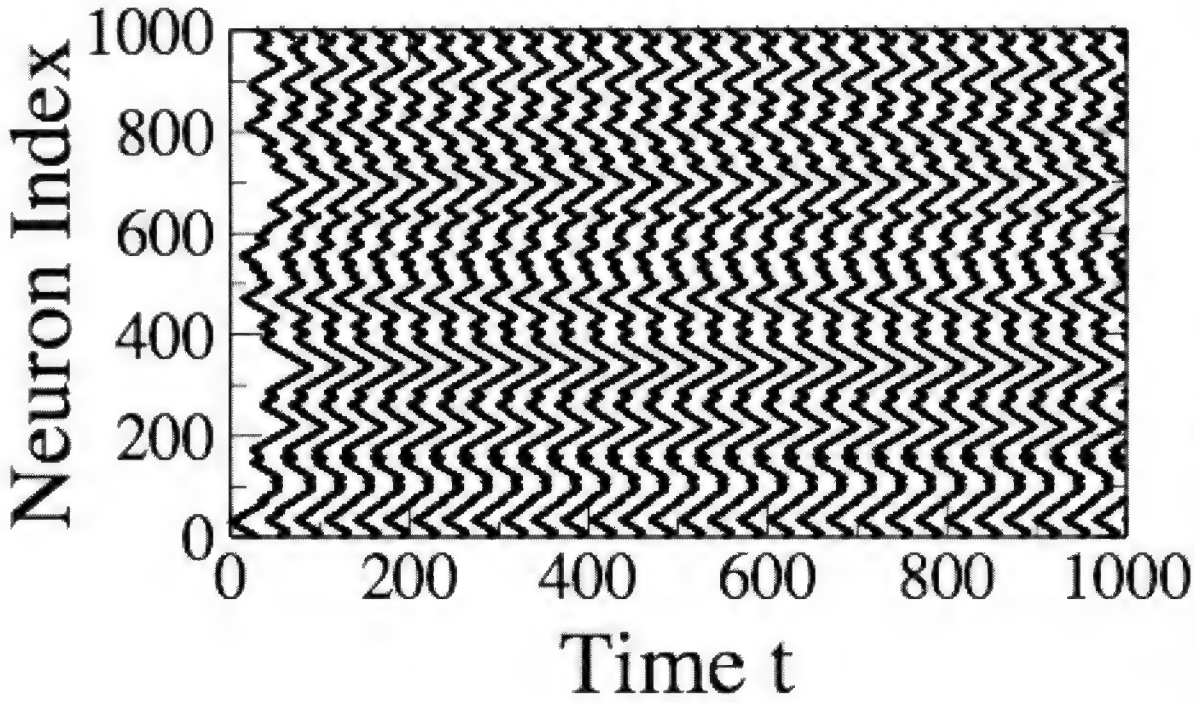

Sustained Network Activity: *p=0.15*

Sustained Network Activity

p=0.05

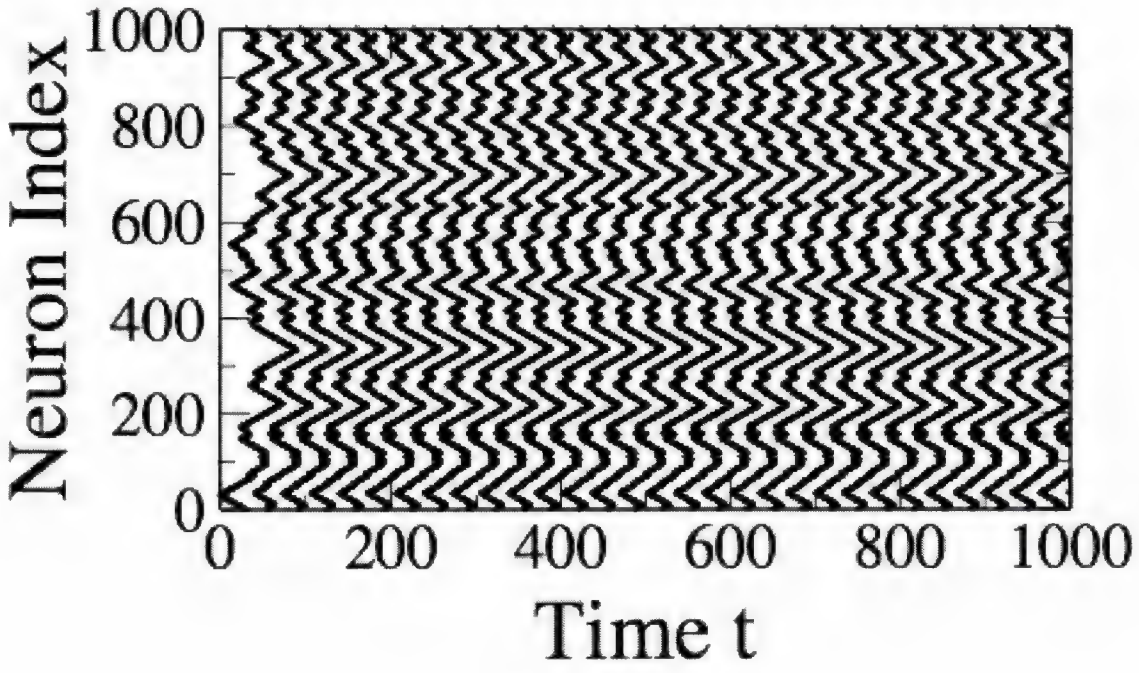

p=0.15

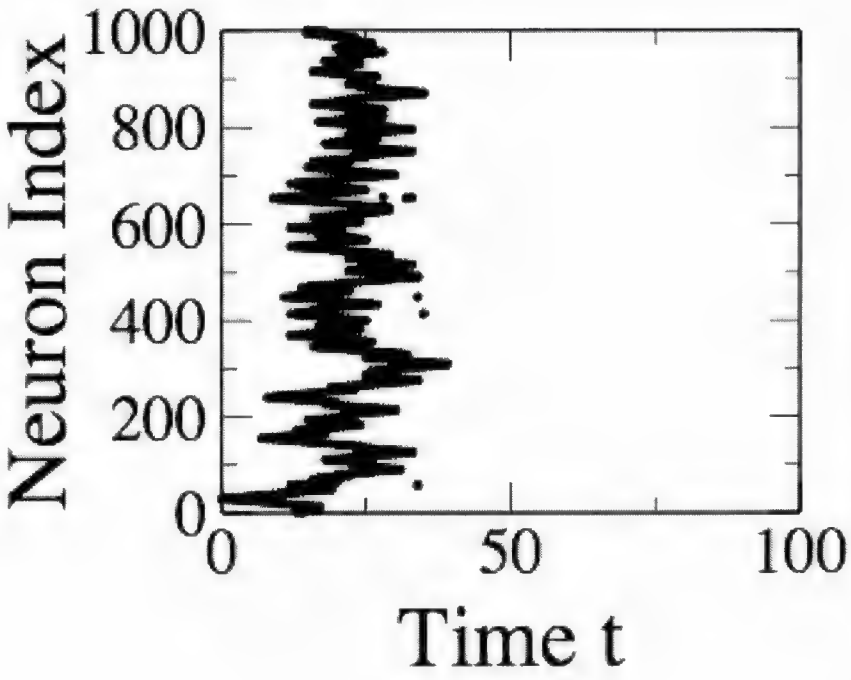

Instantaneous Firing Rate, *p=0.10*

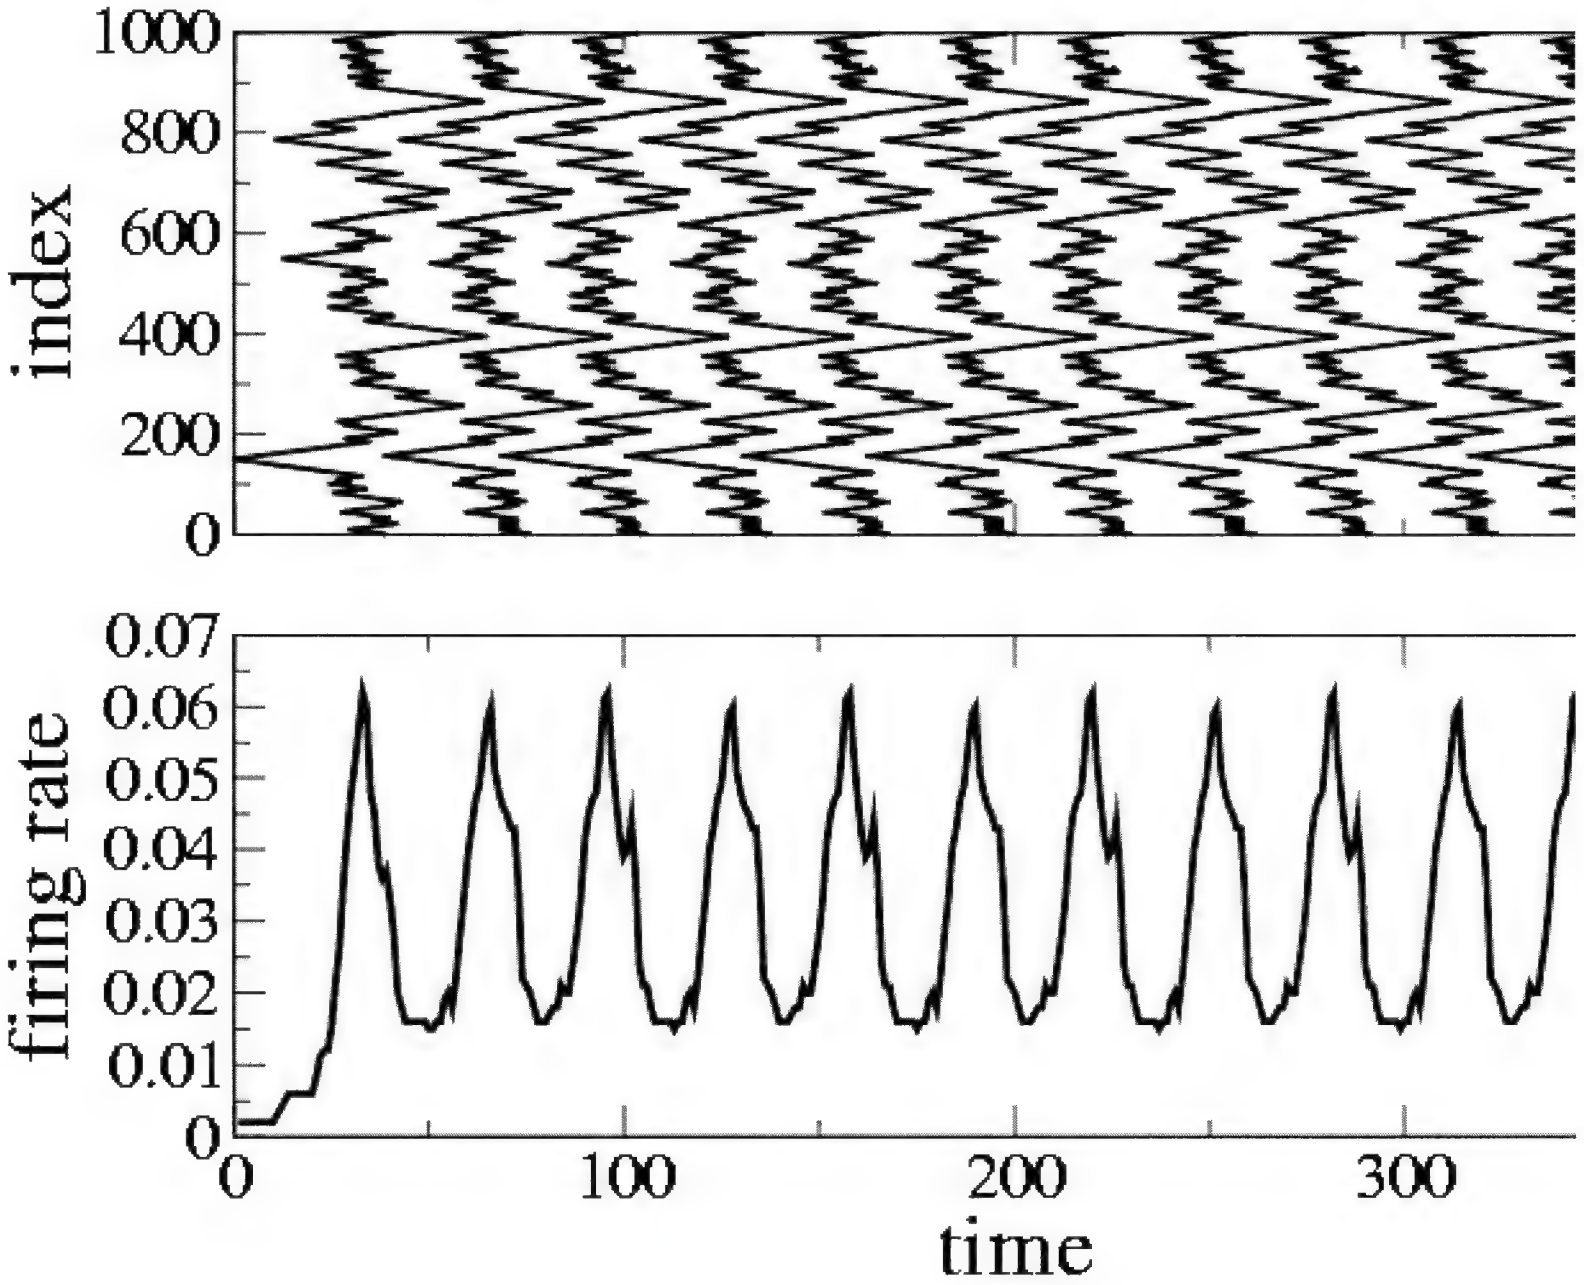

Sustained Network Oscillations

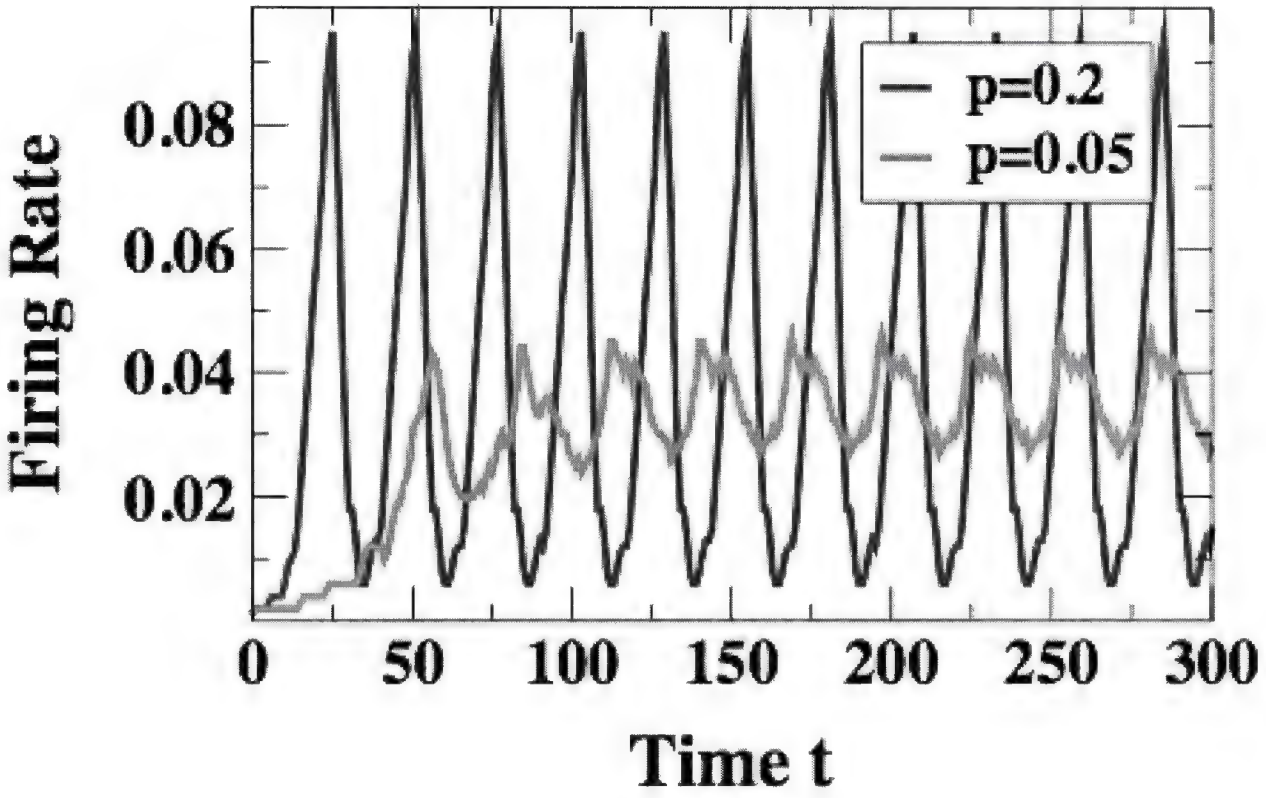

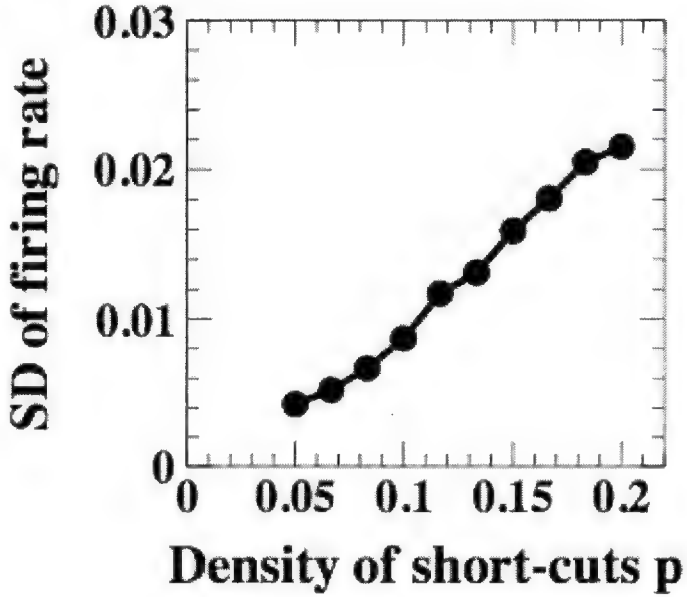

Failure to Sustain Activitv

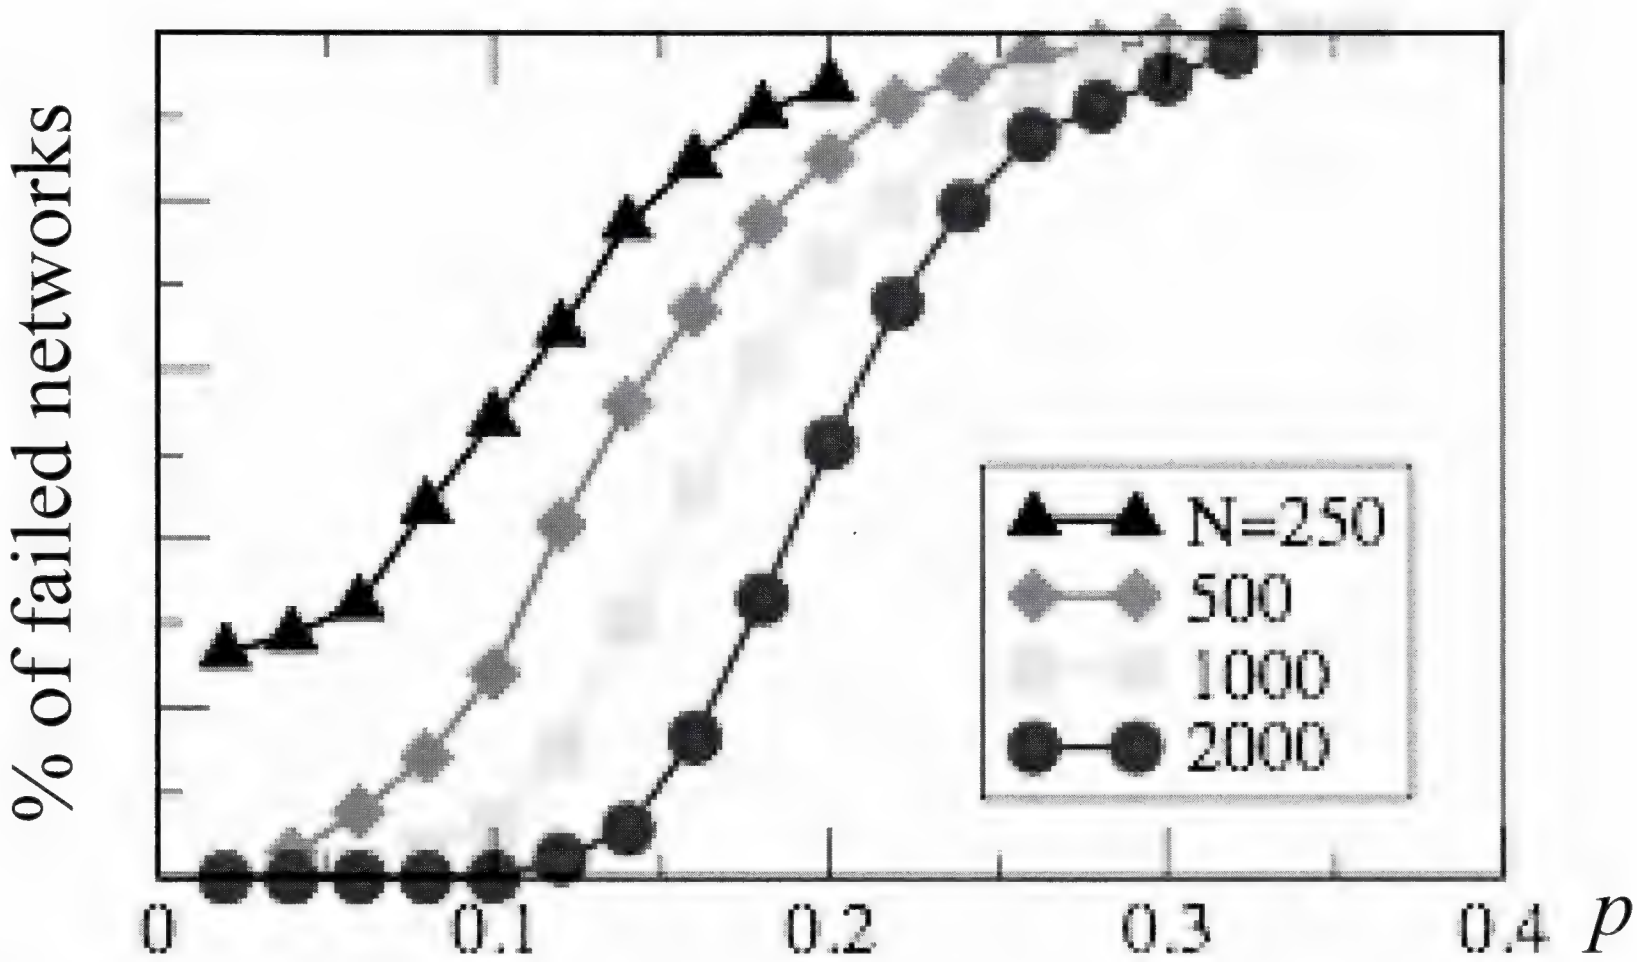

Ensemble average over many network configurations with the same density ***p*** of shortcuts. Some of the configurations will sustain persistent oscillatory activity, while some will burst and fail. Is there a well defined transition for large networks?

Transition to Failure

Failure involves the interaction of two time scales:

1) A cellular time scale associated with the time needed for a neuron to recover to the point where a single synaptic input will make it fire:

$$T_R = \tau \ln\left(\frac{V_\infty - V_{reset}}{V_\infty + V_{syn} - V_{th}}\right), \text{ where } V_\infty - V_{reset} = I/g_L$$

2) A network time scale associated with the first return of activity in a small-world network:

$$\sqrt{\left(1+\frac{4}{pN}\right)}\tanh\left[\sqrt{\left(1+\frac{4}{pN}\right)}\frac{pT_N(p)}{2\tau_D}\right] = 1$$

Newman, Moore, Watts - Phys. Rev. Lett. **84**,

Transition to Failure

The failure transition occurs at a size-dependent critical density of shortcuts, determined by the condition

$$T_N(p_{cr}) = T_R$$

This critical density scales with the logarithm of the size N of the system.

Failure to Sustain Activity With Increasing *p*

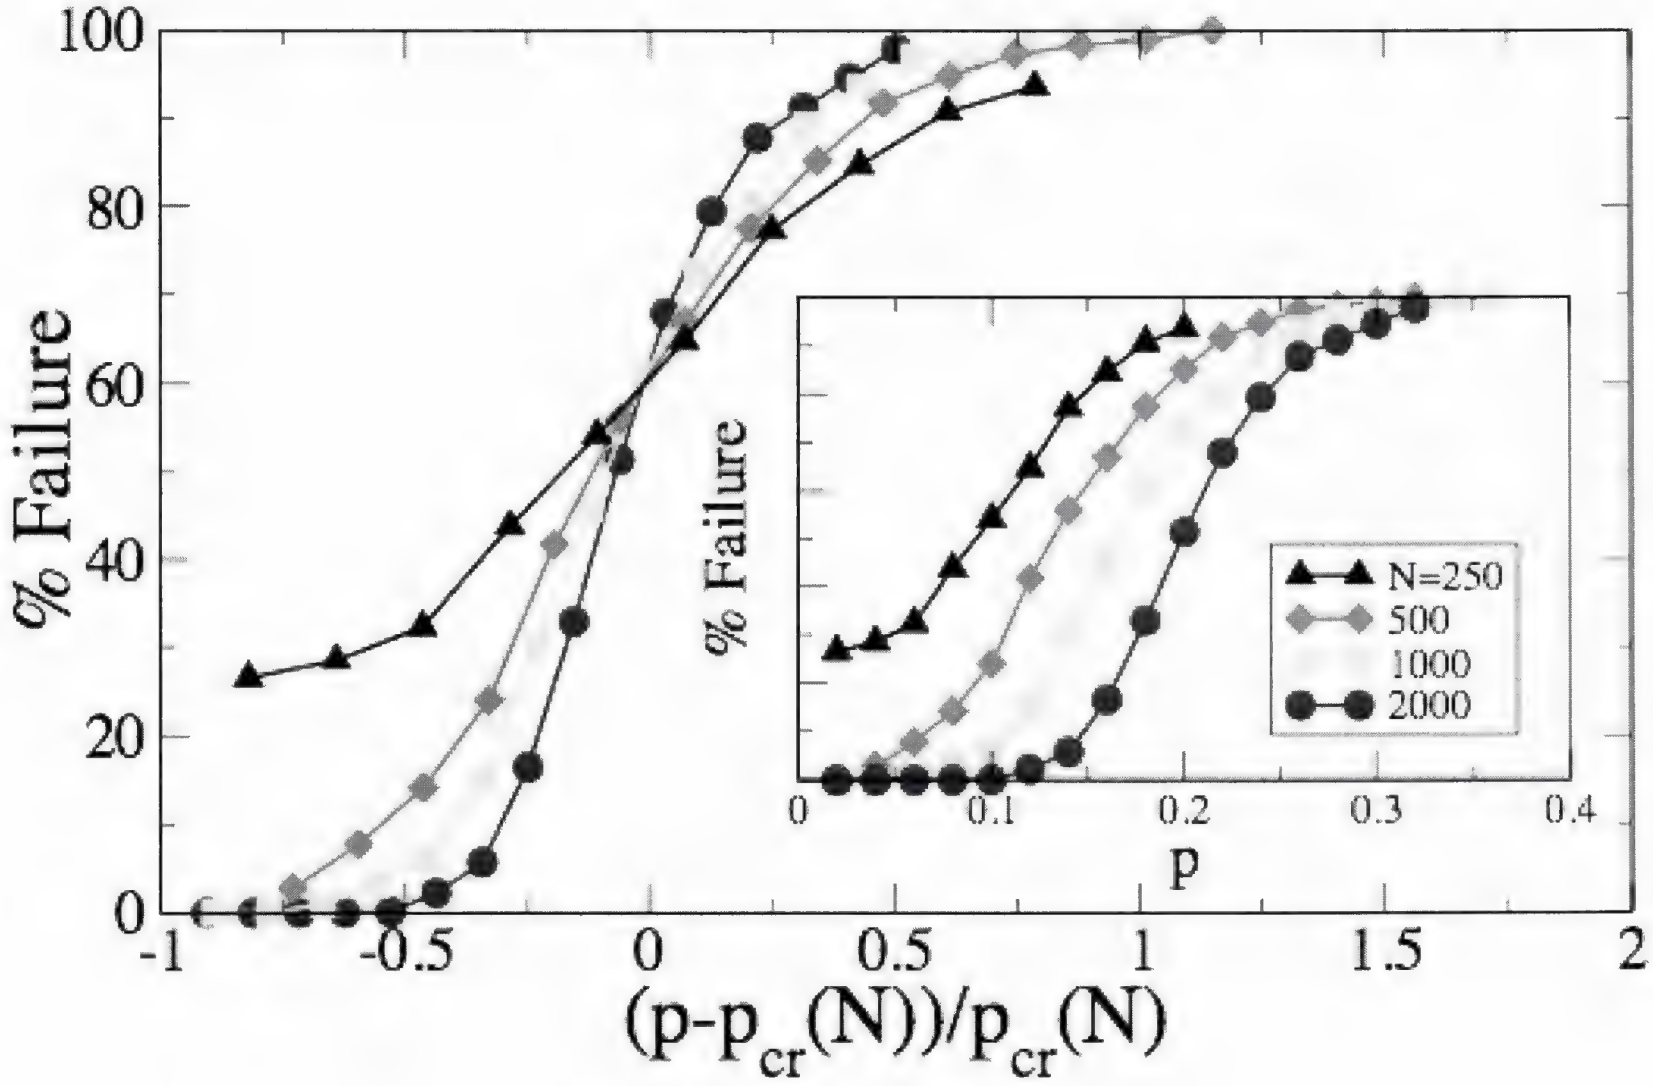

Fast Waves: $\tau_D = 0.10$

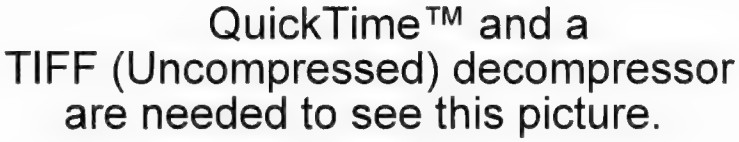

$p = 0.01, 0.05, 0.10, 0.15, 0.20, 0.25$, from (a) to (f)

Sustained Oscillations: Backbone Pathway

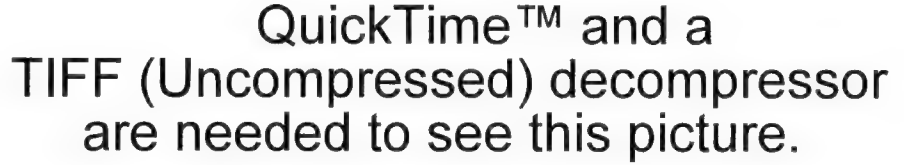

Backbone neurons shown in red; N=1000, p=0.10, and T_R = 2.494.
[Belousov-Zhabotinsky reaction in photosensitive system]

Bistability: a Switch-Off Mechanism

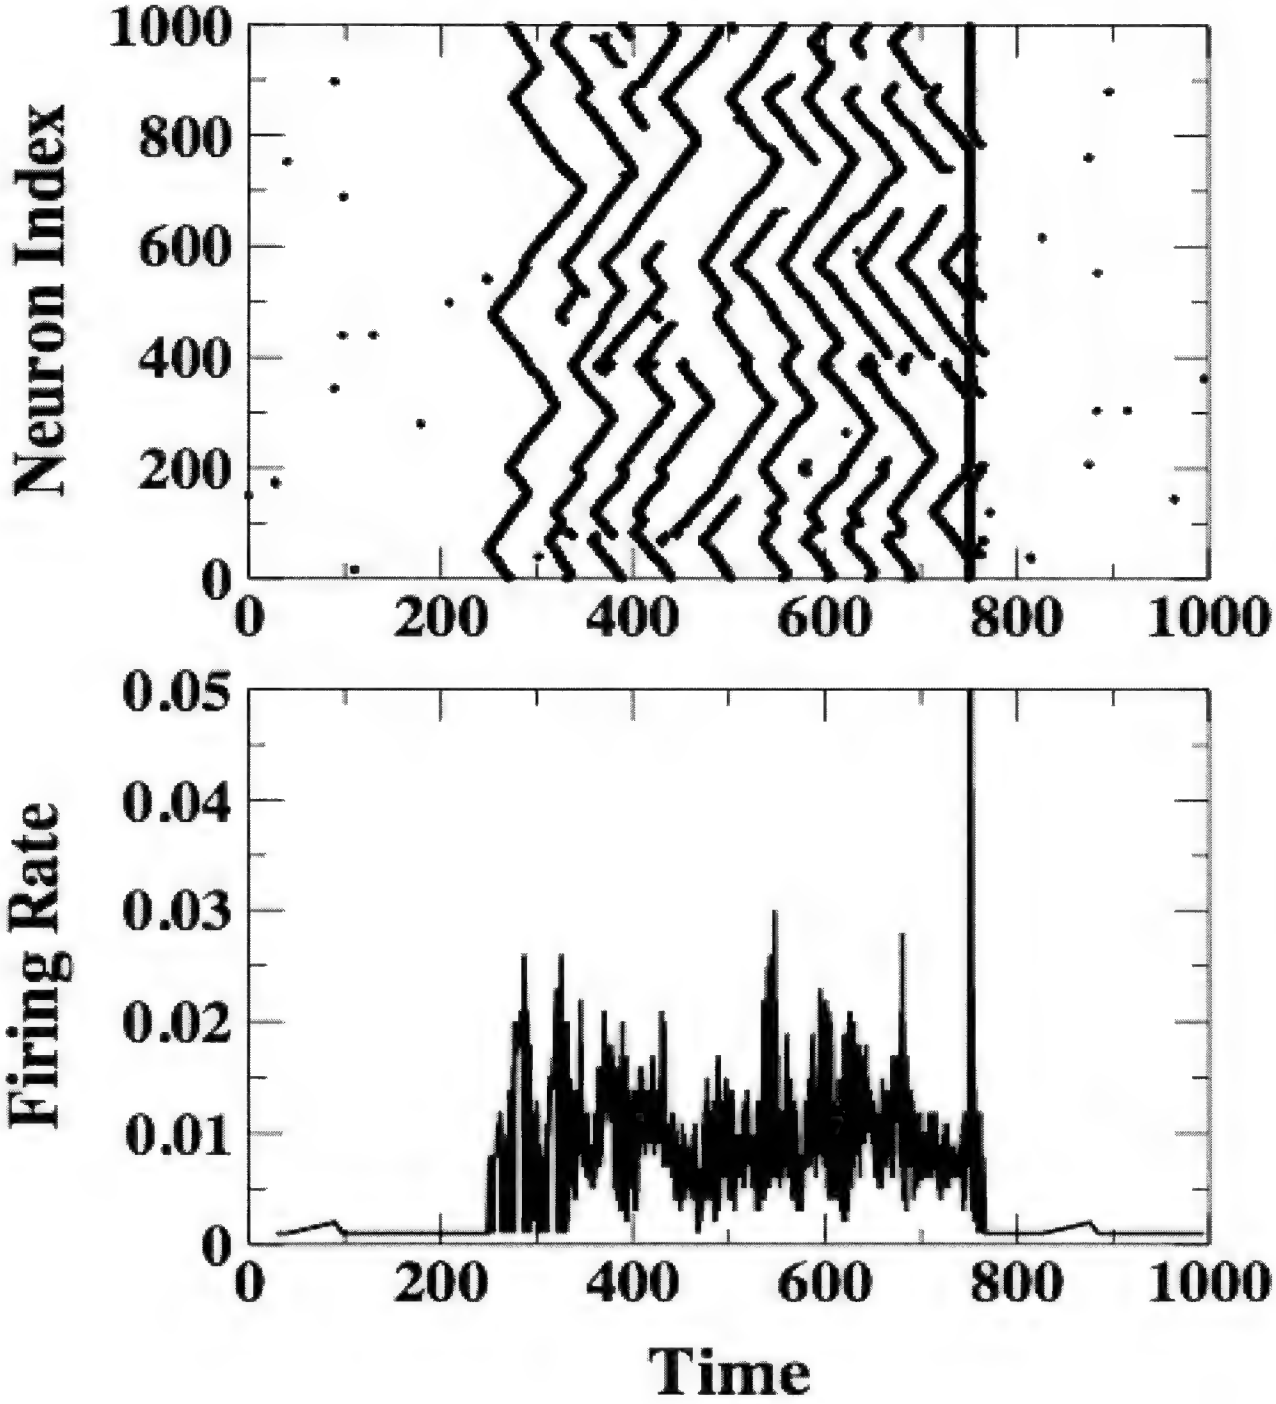

Slow Waves: Increasing τ_D

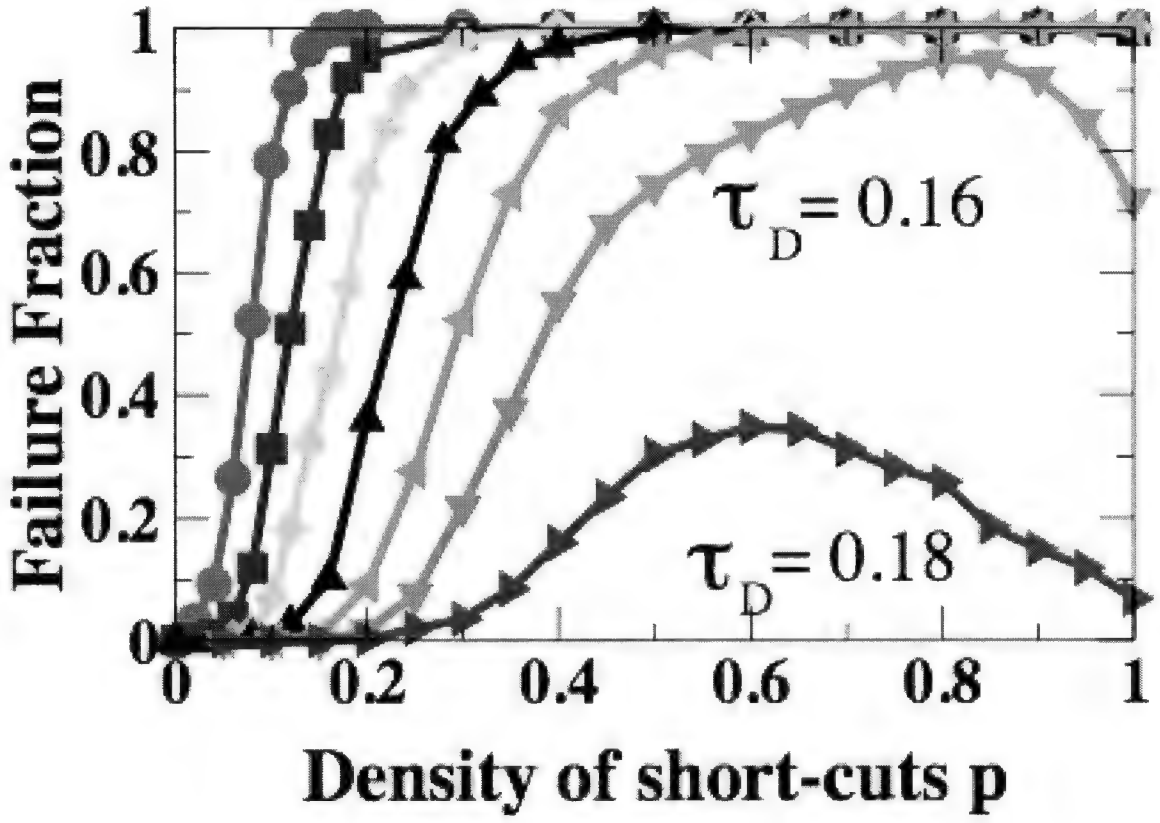

Reentrant Network Activity: $p=0.8$

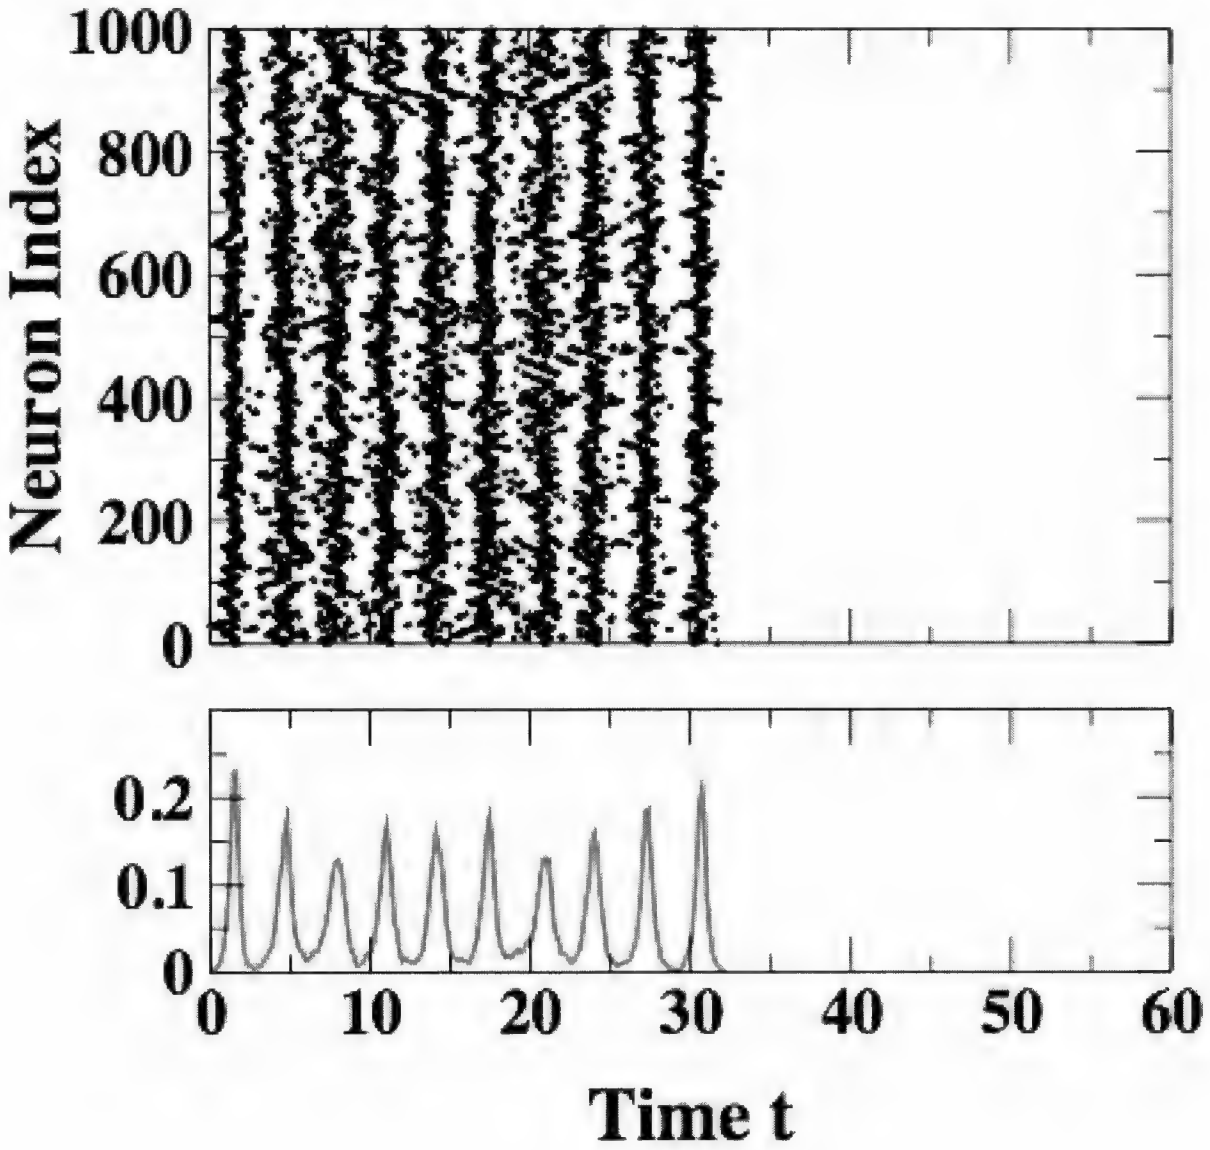

$\tau_D = 0.16$

Reentrant Network Activity: *p=1.0*

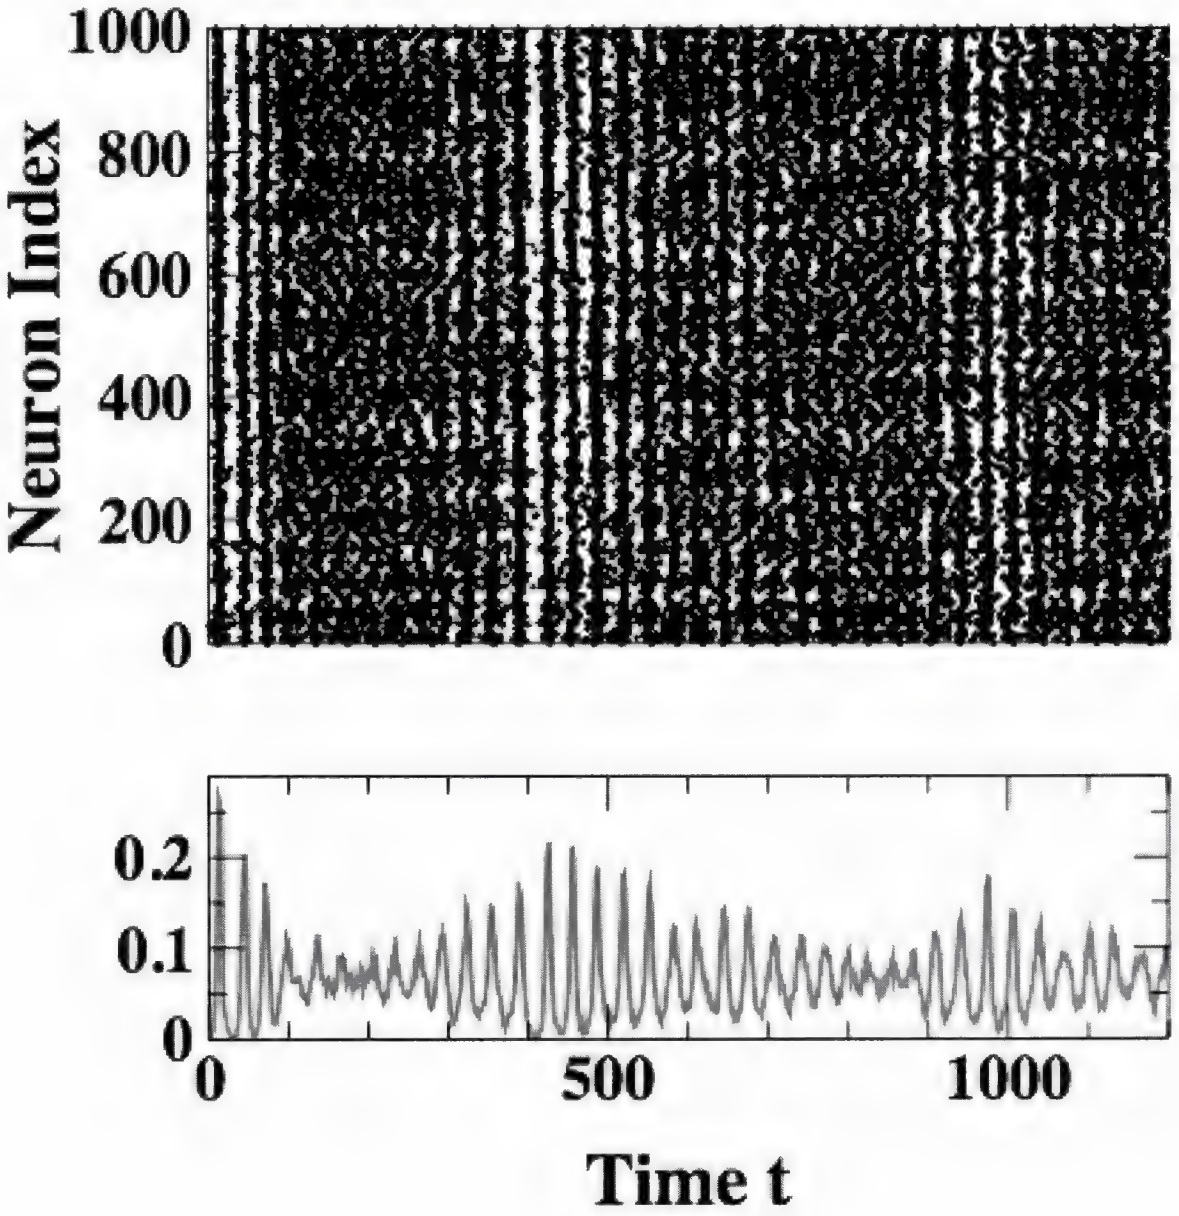

$\tau_D = 0.16$

Slow Waves: $\tau_D = 0.16$

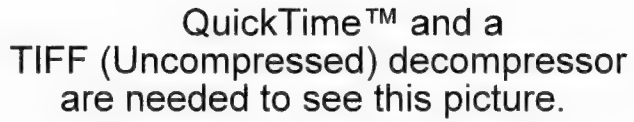

$p = 0.01, 0.20, 0.40, 1.00$, from (a) to (d)

Activity in (d) is noisy and exhibits synchronized population spikes

Chaotic Neural Activity

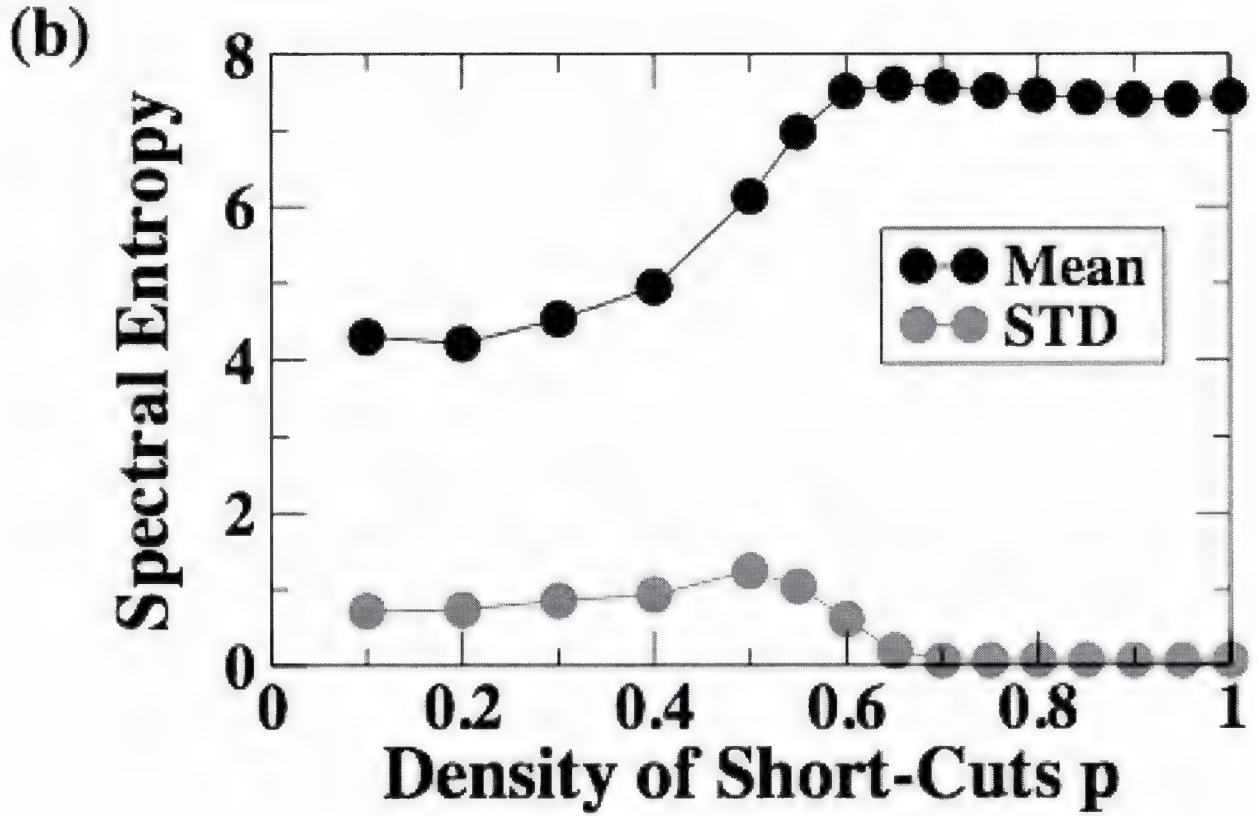

$\tau_D = 0.18$

Properties of the Chaotic Regime

• Below failure transition, number of attractors increases at least linearly with the size N of the system.

• For each attractor, a connectivity backbone can be identified.

• In the chaotic regime, periods of low activity are mediated by *early firing* neurons. These neurons receive more than on shortcut input.

• In the chaotic regime, activity eventually decays. The failure times follow a stretched exponentially distribution.

Summary

- Small-world networks of excitable neurons are capable of supporting sustained activity. This activity is sparse and oscillatory, and it does not require excitatory-inhibitory interactions.

- A transition to failure occurs with increasing density of shortcuts.

- Above this transition, the network dynamics exhibit exceedingly long chaotic transients. This chaotic activity does not require a balanced excitatory-inhibitory network.

Epilepsy: a Network Disorder?

What is an epileptic seizure: sudden, highly synchronized neural activity, which may occur in virtually every area of the brain, and which disrupts normal brain function. What triggers spontaneous seizures?

NETWORK? A leading hypothesis points to axonal sprouting, the growth of axons that do not contact the next relay in the signal pathway but instead return to their original neural network. Sprouted axons provide a pathway by which action potentials can return to and spread through their network of origin, thereby providing the positive feedback needed to sustain seizure activity. [Wuarin, Dudek - J. Neurophysiol. **85**, 1067 (2001)]

Dentate Gyrus: a Model

In temporal lobe epilepsy, changes in synaptic and intrinsic properties occur on a background of altered network architecture resulting from cell loss and axonal sprouting. Question: do structural changes that actually take place during epileptogenesis result in hyperexcitability?

Approach: a data-driven model that incorporates key structural and functional properties of the biological network. Outcome: a network model that incorporates type–specific connectivity data for the rat dentate gyrus. A 1:1 structural model has slightly over 1,000,000 cells, as many as in the dentate gyrus. It incorporates all eight types of dentate cells, with the correct frequency of each type and the correct degree of connectivity.

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**, 1566 (2007)]

Dentate Gyrus: Connectivity Matrix

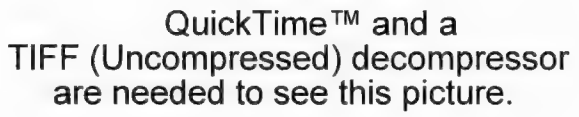

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**,

1:1 Structural Model of the Dentate Gyrus

GC, granule cells;
BC, basket cells;
AAC, axo-axonic cells;
MOPP, molecular layer interneurons with axons in perforant-path termination zone;
MC, mossy cells;
HIPP, hilar interneurons with axons in perforant-path termination zone;
HICAP, hilar interneurons with axons in the commissural/ associational pathway termination zone;
IS, interneuron selective cells.

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**,

Epilepsy in the Dentate Gyrus

An original epileptic event is usually followed by epileptogenesis: recurrent spontaneous seizures. This process is accompanied by 'end-folium sclerosis', a structural reorganization of the dentate gyrus network. The two key components of these morphological changes are the loss of hilar cells and the sprouting of mossy fibers among granule cells.

Approach: Functional model of the dentate gyrus incorporates 50,000 biophysically realistic, multicompartmental single-cell models of excitatory and inhibitory neurons connected by weighted synapses. The model retains four types of cells: the two excitatory cells (i.e., the granule cells and the mossy cells) and two types of interneurons (the somatically projecting fast spiking IS basket cell and the dendritically projecting HIPP cells. Structural changes are implemented by removing mossy cells and by adding mossy fiber contacts among granule cells.

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol.

1:1 Structural Model of the Dentate Gyrus

GC, granule cells;
BC, basket cells;
AAC, axo-axonic cells;
MOPP, molecular layer interneurons with axons in perforant-path termination zone;
MC, mossy cells;
HIPP, hilar interneurons with axons in perforant-path termination zone;
HICAP, hilar interneurons with axons in the commissural/ associational pathway termination zone;
IS, interneuron selective cells.

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**,

1:20 Functional Model of the Dentate Gyrus

Repeated seizures cause loss of hilar cells (both mossy cells and hilar interneurons) and mossy fiber (granule cell axon) sprouting.

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Small-world connectivity is preserved: the predominantly local mossy fiber sprouting compensates for the loss of hilar cells as long as some of the long-distance projecting hilar cells survive.

[Dyhrfjeld-Johnsen, Santhakumar, Morgan, Huerta, Tsimring, Soltesz - J. Neurophysiol. **97**,

Sclerotic Dentate Gyrus

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Changes in network topology predict the development of hyperexcitability. Both normal and low-sclerosis networks exhibit connectivity above p_{cr} and can only sustain bursts. As p decreases below p_{cr}, activity switches to sustained oscillations. The survival of even a few mossy cells is sufficient for maintaining small-world connectivity and thus excitability in the epileptic dentate gyrus.

CONNECTIVITY MATTERS!!!

Self-Sustained Activity in a Small-World Network of Excitable Neurons

Alex Roxin,[1,*] Hermann Riecke,[1] and Sara A. Solla[2]

[1]*Engineering Science and Applied Mathematics, Northwestern University, Evanston, Illinois 60208, USA*

[2]*Department of Physiology and Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208, USA*

(Received 9 June 2003; revised manuscript received 17 February 2004; published 11 May 2004)

We study the dynamics of excitable integrate-and-fire neurons in a small-world network. At low densities p of directed random connections, a localized transient stimulus results either in self-sustained persistent activity or in a brief transient followed by failure. Averages over the quenched ensemble reveal that the probability of failure changes from 0 to 1 over a narrow range in p; this failure transition can be described analytically through an extension of an existing mean-field result. Exceedingly long transients emerge at higher densities p; their activity patterns are disordered, in contrast to the mostly periodic persistent patterns observed at low p. The times at which such patterns die out follow a stretched-exponential distribution, which depends sensitively on the propagation velocity of the excitation.

DOI: 10.1103/PhysRevLett.92.198101 PACS numbers: 87.18.Sn, 82.40.Bj, 89.75.Hc

Recent research in complex networks has provided increasing evidence for their relevance to a variety of physical, biological, and social phenomena [1–3]. Two distinct types of topology have been particularly useful in providing insights into the implications of complex connectivity: scale-free networks [3], characterized by the existence of a small number of hubs with high coordination number, and small-world networks [1], characterized by the presence of shortcuts that link two randomly chosen sites regardless of the distance between them.

So far, most work on complex networks has focused on their topological and geometrical properties; less attention has been given to the properties of dynamical systems defined on such networks. The interplay between the intrinsic dynamics of the constituent elements and their complex pattern of connectivity strongly affects the collective dynamics of the resulting system. For instance, the addition of shortcuts induces a finite-temperature phase transition even in the one-dimensional Ising model [4], and the introduction of unidirectional shortcuts can change the second-order phase transition in the two-dimensional Ising model into a first-order one [5]. In a system of coupled oscillatory elements, the introduction of shortcuts enhances synchronization [6], while the introduction of hubs eliminates the threshold for epidemic propagation [7].

The coexistence of shortcuts and regular local connections characteristic of small-world networks (SWNs) mimics a salient feature of the circuitry in the cortex [8–13], where experimental observations of excitatory traveling waves [11] provide evidence of some degree of local connectivity, while it is also recognized that long-range excitatory connections are present [10,12]. Our goal is to explore the influence of this complex connectivity on the dynamics of neuronal circuits; to this purpose, we choose a minimal model. The underlying network is modeled as a SWN with unidirectional shortcuts that reflect the nonreciprocal character of synaptic connections, and the excitable neurons are modeled as leaky integrate-and-fire units. We find that even this simple model exhibits a rich repertoire of distinct dynamical behaviors as a function of the density p of added shortcuts: a low p regime characterized by persistent periodic activity that is bistable with the quiescent state, a transition to failure with increasing p, followed by a reemergence of long-lasting disordered activity. We note that a SWN with unidirectional shortcuts has been considered in a different regime by Lago-Fernández *et al.* [14,15] to address the possibility of rapid synchronization among conductance-based neurons of the Hodgkin-Huxley type.

The model considered here consists of a one-dimensional array of N integrate-and-fire neurons (IFNs) in which a SWN topology is created through the addition of a density p of unidirectional long-range couplings. The membrane potential of the IFNs is determined by

$$\tau_m \frac{dV_i}{dt} = -V_i + I_{\text{ext}} + g_{\text{syn}} \sum_{j,m} w_{ij} \delta(t - t_j^{(m)} - \tau_D). \quad (1)$$

A neuron fires whenever its voltage exceeds 1; the voltage is then reset to 0. The external current is chosen to satisfy $I_{\text{ext}} < 1$; in this regime the IFNs are not oscillatory, but excitable. The last term models input currents due to presynaptic firing as a delayed impulse: if w_{ij} is nonzero, then neuron i receives a pulse input of amplitude g_{syn} with a delay τ_D after neuron j has fired its mth spike at time $t_j^{(m)}$. The synaptic conductance is chosen to satisfy $I_{\text{ext}} + g_{\text{syn}} > 1$, so that a single input suffices to sustain firing activity. The local connections are modeled here as nearest-neighbor couplings ($w_{i,i\pm1} = 1$) that define an underlying regular lattice. The long-range connections result from randomly *adding* rather than rerouting [1] a fixed fraction pN of unidirectional couplings $w_{ij} = 1$ to generate a SWN topology.

At $p = 0$, any excitation sufficient to cause a neuron to fire will generate two pulses that propagate through the

 © 2004 The American Physical Society

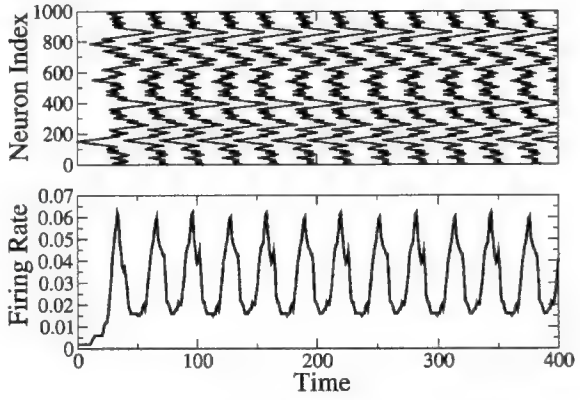

FIG. 1. Raster plot (top) and instantaneous firing rate (bottom) for a system with $N = 1000$, $I_{\text{ext}} = 0.85$, $g_{\text{syn}} = 0.2$, $\tau_m = 10$, $\tau_D = 1$, $p = 0.1$. Same parameter values are used in subsequent figures unless noted otherwise.

regular lattice in opposite directions with velocity $v = 1/\tau_D$, and either exit the system or annihilate each other, depending on boundary conditions. No persistent activity results in either case. However, self-sustained activity may arise for nonzero p, as shown in Fig. 1. Persistence relies on the reinjection of activity via a shortcut into a previously active domain that has by then recovered; this reinjection can occur only if the shortcuts are unidirectional. For a fixed value of p, any particular network realization has a different connectivity graph that may or may not sustain persistent activity. We typically average over 2000 realizations to calculate the probability of persistent activity; the complementary probability of *failure* to sustain activity is shown in Fig. 2 (inset) as a function of the density p of random connections for four different system sizes. In this regime, the probability of failure is an increasing function of p that crosses over from 0 to 1 with increasing steepness as the size N of the system increases.

Failure to sustain activity is a simple consequence of the intrinsic dynamics of the neurons. Pulses travel outwards from an initial activity seed and spawn new pulses at a rate that increases with p. A currently inactive neuron can fire again only if enough time has elapsed from its preceding firing to allow for a recovery to $V \geq 1 - g_{\text{syn}}$. A single input will be able to elicit a spike only if the elapsed time exceeds $T_R^{(1)}$, with

$$T_R^{(n)} = \tau_m \ln\left(\frac{I_{\text{ext}}}{I_{\text{ext}} + n g_{\text{syn}} - 1}\right). \tag{2}$$

If activity recurs to a given site too rapidly, the neuron will fail to produce a spike, and the pulse of activity will die out. A critical density p_{cr} for the transition from self-sustained activity to failure can be estimated from

$$T_A(p_{cr}) = T_R^{(1)}, \tag{3}$$

where $T_A(p)$ is the time needed for the activity to spread

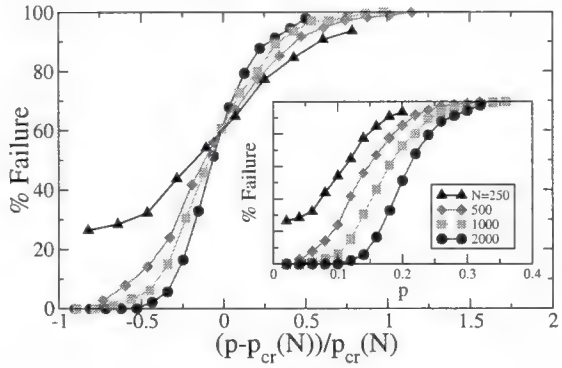

FIG. 2 (color online). Inset: Failure rates for $\tau_D = 1.0$ and different system sizes N. Main: Failure rates normalized by $p_{cr}(N)$ [cf. (3) and (4)].

across the whole network. At a fixed velocity for pulse propagation, this time corresponds to the largest distance across the network. This distance has been calculated for bidirectional shortcuts using a mean-field approach [16]; when extended to the case of unidirectional shortcuts it results in

$$\sqrt{\left(1 + \frac{4}{pN}\right)} \tanh\left[\sqrt{\left(1 + \frac{4}{pN}\right)} \frac{pT_A(p)}{2\tau_D}\right] = 1. \tag{4}$$

An implicit expression for p_{cr} and its dependence on the system size N and the propagation velocity v follows from combining Eqs. (3) and (4); for large N, $p_{cr}(N) \propto \ln N$. Failure rate curves as a function of $[p - p_{cr}(N)]/p_{cr}(N)$, shown in Fig. 2, cross at the theoretically predicted value. This observation, together with the increased steepness of these curves with increasing N, indicates that a well-defined transition to failure occurs in the thermodynamic limit.

This well-defined transition to failure occurs only for sufficiently fast waves, i.e., for short delay τ_D. For larger τ_D the dynamics of the system become quite more complex, and the fraction of realizations that fail before a fixed time ($t_{\text{max}} = 2000$ in Fig. 3) becomes a nonmonotonic function of p. While at low p the firing patterns are highly regular (cf. Fig. 1) and all failures occur within one or two cycles of the initial activity, for higher p the patterns are more disordered (cf. Fig. 4) and the activity can persist for a very long time before failure.

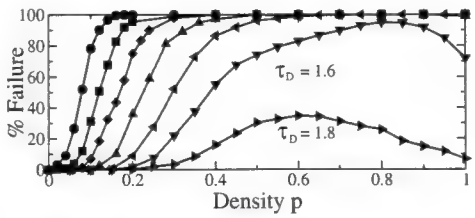

FIG. 3. Failure rates after $t_{\text{max}} = 2000$ for $\tau_D = 0.6$, 0.8, 1.0, 1.2, 1.4, 1.6, 1.8 (left to right); $N = 1000$.

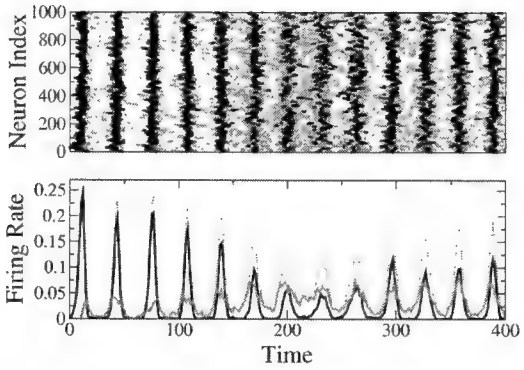

FIG. 4 (color online). Raster plot (top) and instantaneous firing rate (bottom) of neurons with ISI $> T_R^{(1)}$ (black) and ISI $< T_R^{(1)}$ [gray (red online)] for $\tau_D = 1.5$ and $p = 1.0$. Dotted line is the total firing rate.

Consequently, the distribution of failure times, shown in Fig. 5, exhibits an increasingly long tail for longer delay times.

To understand the persistence of activity beyond p_{cr}, it is important to recognize that the result for the critical density $p_{cr}(N)$ hinges on the assumption that each neuron receives but a single excitatory input during each cycle of network activity. Its recovery time is therefore given by $T_R^{(1)}$, which sets a lower bound for the interspike interval (ISI). While this assumption is well satisfied for small p, it does not hold for $p = O(1)$. In fact, the likelihood that a neuron has n incoming shortcuts follows a multinomial distribution such that the fraction of neurons receiving two incoming connections grows from about 0.05% at $p = 0.1$ to about 30% at $p = 1$. Neurons that receive n inputs can have a recovery time as short as $T_R^{(n)}$. Figure 4 reveals that such neurons, with ISIs lower than $T_R^{(1)}$, play a crucial role in maintaining network activity in the regime $p \approx 1$. While neurons with ISIs greater than $T_R^{(1)}$ (shown black in Fig. 4) can go through silent epochs with near-zero activity, neurons with shorter ISIs [gray (red online) in Fig. 4] may fire 2 or 3 times during a network cycle and carry over the network activity across these silent epochs.

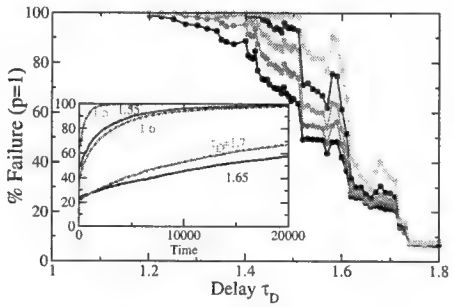

FIG. 5 (color online). Failure-time distribution. Inset: Cumulative distribution of failure times at $p = 1$ for $1.5 \leq \tau_D \leq 1.7$. Main: Failure rates at 5, 10, 20, 40, and 100 multiples of $T_R^{(1)} = 28.3$ (bottom to top). Symbols are averages over 2000 realizations.

The persistence of activity beyond p_{cr} depends sensitively on τ_D; this reflects the fact that persistence at higher densities p is due to chains or trees of neurons that bridge the silent epochs due to their multiple inputs. Since activity propagates with a fixed speed, a chain of multiple-input neurons of given length can bridge a time interval proportional to τ_D. Thus, as τ_D increases, ever shorter chains can contribute to bridging a silent epoch of given duration; the likelihood for failure will decrease accordingly. This picture is, of course, overly simplistic: whether a *topological* chain can be utilized as a *dynamical* bridge over a given silent epoch depends on the amount and timing of the inputs it receives, which in turn depend on the recent history of the entire network. Simulations reveal that the identity of the neurons that form the "bridging" dynamical chains varies from cycle to cycle in an irregular way. This implies that the effective utilization of a dynamical chain over one cycle does not guarantee its availability on the next cycle. Therefore, even systems that persist for long times may still have a finite probability of failing.

The cumulative failure distribution function $F(t)$, shown in Fig. 5 for various values of τ_D, exhibits a long tail and is well fit with stretched exponentials: $F(t) = f_\infty(\tau_D) - Ce^{-\alpha t^\beta}$ with $\beta \approx 0.4$ (dotted lines). Even though the fits are based on runs up to $t = 300\,000$ for $\tau_D = 1.65$ and $\tau_D = 1.7$, they do not provide a value of $f_\infty(\tau_D)$ accurate enough for establishing whether true persistent activity exists for a small fraction of the network realizations [$0.97 \leq f_\infty(\tau_D = 1.65) \leq 1$]. Strikingly, the dependence of $F(t)$ on the delay time exhibits a high degree of structure, suggestive of "resonances" at values of τ_D for which some chains and trees can be optimally utilized.

One of the salient features of the emergent dynamics of the model is persistent *self-sustained* activity at low densities p of shortcuts. In this regime, the network is, in fact, bistable between "off" and "on" states, and can be switched between them with sufficiently large stimuli, as illustrated in Fig. 6. The synchronous stimulation of a sufficiently large number of neurons while the network is in the on state increases the level of activity and effectively pushes the network to the right of the failure transition (cf. Fig. 2), causing a transition into the off state.

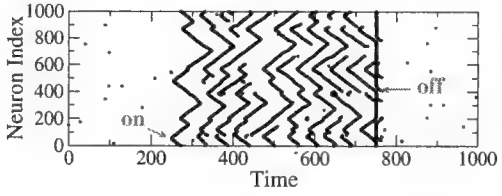

FIG. 6 (color online). Raster plot for $p = 0.10$ and $k = 5$. About 10 adjacent neurons are stimulated synchronously at $t = 250$; about 20% of the neurons are activated at $t = 750$.

To address the fact that real neuronal systems are noisy, the simulation shown in Fig. 6 includes Gaussian fluctuations in the membrane potential; their amplitude is chosen so as to cause the neurons to spike irregularly at a low rate. To keep the noisy spiking of a single neuron from generating a traveling pulse and initiating the on state, we adjust the synaptic conductance so that several adjacent neurons must fire in rapid succession in order to propagate a pulse of excitation. The network topology is modified accordingly: we extend the local coupling to include up to $2k$ neighbors ($w_{i,i\pm j} = 1$ for $j = 1, \ldots, k$) and model long-range connections via a population of intermediate excitatory neurons that both receive input from and project to multiple adjacent neurons. Under these conditions, spontaneous activity is highly unlikely to initiate traveling pulses. However, a sufficiently large stimulus, synchronous across several neurons, can again turn the state of elevated activity on and off (cf. Fig. 6); bistability is thus robust with respect to noise.

Network bistability has been hypothesized to be the neural correlate underlying the type of short-term memory known as working memory in the prefrontal cortex of monkeys and humans. Much more realistic and physiologically plausible models of cortical layers have been studied within the context of working memory (e.g., [17]). Yet, not much attention has been given to heterogeneities in network topology or to long-range excitatory connections. The work presented here suggests that closer attention be given to the role of connectivity as an additional factor that contributes to the generation of the persistent, active state associated with working memory.

In conclusion, we have investigated the effect of incorporating random unidirectional shortcuts to a one-dimensional network of locally coupled integrate-and-fire neurons. We find that even a very low density of shortcuts suffices to generate persistent activity from a local stimulus through the reinjection of activity into previously excited domains. As the density of shortcuts is increased, the substantial decrease in the effective system size characteristic of small-world networks causes a crossover into a regime characterized by failure to sustain activity for essentially all network configurations. For sufficiently slow propagation velocities of the activity and sufficiently high shortcut densities, an intriguing second crossover occurs into a regime in which the activity still fails but only after often exceedingly long and strongly chaotic transients.

The complex dynamical phenomena we find in this extremely simple model are based on a robust mechanism: propagating pulses of activity that are sustained by branching and reinjection. We therefore expect that more realistic models of neuronal networks, which may include multiple ion channels and continuous synaptic currents as well as inhibitory coupling, will show qualitatively the same behavior upon the addition of shortcuts if they originally support propagating pulses of activity [18] that annihilate upon collision. Preliminary computations show that the failure transition persists if the length of the shortcuts is limited by $L_{\max} > T_R^{(1)}/\tau_D$ [19]; this network is not truly a SWN. If the SWN is obtained by rerouting rather than adding connections, the transition is less pronounced [19]. Recently, similar networks have been used to simulate epileptic activity in hippocampus [20] and bursting in the pre-Bötzinger complex [21]. The phenomena reported here should also be accessible in excitable chemical systems [22], where shortcuts could be implemented using video feedback.

We gratefully acknowledge support by DOE (DE-FG02-92ER14303), by NSF Grant No. DMS-9804673, and by the NSF-IGERT program *Dynamics of Complex Systems in Science and Engineering* (DGE-9987577), as well as discussions with R. Clewley and R. J. Butera.

*Present address: Neurophysics and Physiology of the Motor System, Université René Descartes, Paris, France.

[1] D. J. Watts and S. H. Strogatz, Nature (London) **393**, 440 (1998).
[2] S. H. Strogatz, Nature (London) **410**, 268 (2001).
[3] R. Albert and A. L. Barabasi, Rev. Mod. Phys. **74**, 47 (2002).
[4] A. Pekalski, Phys. Rev. E **64**, 057104 (2001).
[5] A. D. Sánchez, J. M. López, and M. A. Rodríguez, Phys. Rev. Lett. **88**, 048701 (2002).
[6] M. Barahona and L. M. Pecora, Phys. Rev. Lett. **89**, 054101 (2002).
[7] R. Pastor-Satorras and A. Vespignani, Phys. Rev. Lett. **86**, 3200 (2001).
[8] V. B. Mountcastle, Brain **120**, 701 (1997).
[9] G. González-Burgos, G. Barrionuevo, and D. A. Lewis, Cereb. Cortex **10**, 82 (2000).
[10] D. A. McCormick and D. Contreras, Annu. Rev. Physiol. **63**, 815 (2001).
[11] J. Y. Wu, L. Guan, and Y. Tsau, J. Neurosci. **19**, 5005 (1999).
[12] R. Demir, L. B. Haberly, and M. B. Jackson, J. Neurophysiol. **86**, 2445 (2001).
[13] D. S. Melchitzky, G. González-Burgos, G. Barrionuevo, and D. A. Lewis, J. Comp. Neurol. **430**, 209 (2001).
[14] L. F. Lago-Fernández, R. Huerta, F. Corbacho, and J. A. Sigüenza, Phys. Rev. Lett. **84**, 2758 (2000).
[15] L. F. Lago-Fernández, F. J. Corbacho, and R. Huerta, Neural Networks **14**, 687 (2001).
[16] M. E. Newman, C. Moore, and D. J. Watts, Phys. Rev. Lett. **84**, 3201 (2000).
[17] A. Compte, N. Brunel, P. S. Goldman-Rakic, and X.-J. Wang, Cereb. Cortex **10**, 910 (2000).
[18] D. Golomb and G. B. Ermentrout, Phys. Rev. Lett. **86**, 4179 (2001).
[19] A. Roxin, H. Riecke, and S. A. Solla (unpublished).
[20] T. I. Netoff, R. Clewley, S. Arno, T. Keck, and J. A. White (to be published).
[21] W. C. Gerken, J. Shao, and R. J. Butera (unpublished).
[22] K. Martinez, A. L. Lin, R. Kharrazian, X. Sailer, and H. L. Swinney, Physica (Amsterdam) **168D**, 1 (2002).

Multiple attractors, long chaotic transients, and failure in small-world networks of excitable neurons

Hermann Riecke
Department of Engineering Sciences and Applied Mathematics, Northwestern University, Evanston, Illinois 60208

Alex Roxin
Computational Neuroscience, Departament de Tecnologia, Universitat Pompeu Fabra, 08003 Barcelona, Spain

Santiago Madruga
Max-Planck-Institute for Physics of Complex Systems, D-01187 Dresden, Germany

Sara A. Solla
Department of Physiology and Department of Physics and Astronomy, Northwestern University, Evanston, Illinois 60208

(Received 3 January 2007; accepted 2 May 2007; published online 28 June 2007)

We study the dynamical states of a small-world network of recurrently coupled excitable neurons, through both numerical and analytical methods. The dynamics of this system depend mostly on both the number of long-range connections or "shortcuts", and the delay associated with neuronal interactions. We find that persistent activity emerges at low density of shortcuts, and that the system undergoes a transition to failure as their density reaches a critical value. The state of persistent activity below this transition consists of multiple stable periodic attractors, whose number increases at least as fast as the number of neurons in the network. At large shortcut density and for long enough delays the network dynamics exhibit exceedingly long chaotic transients, whose failure times follow a stretched exponential distribution. We show that this functional form arises for the ensemble-averaged activity if the failure time for each individual network realization is exponentially distributed. © *2007 American Institute of Physics*. [DOI: 10.1063/1.2743611]

Many systems in nature can be described as a network of interconnected nodes. Networks in a growing list of systems, from social and ecological webs to the neural anatomy of simple organisms, have been shown to exhibit complex topological features that distinguish them from both ordered lattices and purely random networks. Beyond the investigation of the structural and geometrical properties of such networks, a new class of question arises when dynamical degrees of freedom are placed at their nodes. As the investigation of such dynamical complex networks proceeds, it has become increasingly clear that the network architecture can significantly influence the dynamics of the system. An understanding of emergent dynamics on complex networks requires investigating the interplay between the intrinsic dynamics of the node elements and the connectivity of the network in which they are embedded. In order to address some of these questions in a specific scenario of relevance to the dynamical states of neural ensembles, we study here the collective behavior of excitable model neurons in a network with small-world topology. The small-world network has local lattice order, but includes a number of randomly placed connections that may provide connectivity shortcuts. This topology bears a schematic resemblance to the connectivity of the cerebral cortex, in which neurons are most strongly coupled to nearby cells within 50–100 μm, but also make projections to cells millimeters away. We find that the dynamics of this small-world network of excitable neurons depend mostly on both the density of shortcuts and the delay associated with neuronal projections. In the regime of low shortcut density, the system exhibits persistent activity in the form of propagating waves, which annihilate upon collision and are spawned anew via the reinjection of activity through shortcut connections. As the density of shortcuts reaches a critical value, the system undergoes a transition to failure. The critical shortcut density results from matching the time associated with a recurrent path through the network to an intrinsic recovery time of the individual neurons. Furthermore, if the delay associated with neuronal interactions is sufficiently long, activity re-emerges above the critical density of shortcuts. The activity in this regime exhibits long, chaotic transients composed of noisy, large-amplitude population bursts. A numerical investigation of the interplay between network topology and interaction delays in this regime reveals a mechanism that underlies the observed stretched-exponential distribution of failure times for the chaotic network activity.

I. INTRODUCTION

It has been widely recognized that the connectivity of a network of active elements has a profound impact on its function. Substantial effort has therefore been devoted to the characterization of network connectivity,[1,2] leading to the identification of various measures that are significant in de-

 © 2007 American Institute of Physics

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

termining the geometrical properties of the system. Particularly relevant among them are the average and maximal length of the minimal paths that connect two arbitrary nodes in the network, the clustering coefficient, which characterizes the propensity of all neighbors of a given node to be connected to each other, and the distribution for the degree, defined as the number of links that emanate from a node.

A large number of networks with complex topology fall into the class of small-world networks, characterized by short average path length and high clustering coefficient. A simple realization of such a small-world network consists of a regular lattice supplemented by a number of randomly placed connections that tend to provide connectivity shortcuts. This topology schematically resembles the connectivity of the cerebral cortex, in which neurons are most strongly coupled to nearby cells within 50–100 μm, but also project to distant cells, which can be millimeters away. Small-world properties of neuronal networks have been found in a comprehensive morphological characterization of *in vitro* two-dimensional networks,[3] and in a 1:1 network model of the rat dentate gyrus that incorporates known data about cell types, cell-specific connectivity, and axonal branch length.[4]

The dynamics of elements coupled through the connectivity of a complex network has been studied in detail for the case of oscillatory elements. The emphasis has been in determining the role of network topology in the ability to achieve synchronization in a system of coupled oscillators. The existence of long-range connections, which reduce the effective size of the network, has been found to substantially enhance synchronizability.[5,6] At the same time, the heterogeneity of the degree distribution found in many complex networks limits the ability of the oscillators to synchronize. Stable synchronization emerges from a balance between these two competing aspects of complex network topology.[7]

Excitable elements are the components of another important class of dynamical systems, whose emergent behavior is not characterized by synchronization. Locally coupled networks of excitable elements exhibit traveling waves (e.g., Refs. 8–10). If these waves annihilate upon collision, as is typically the case, persistent activity usually requires either an external drive or spontaneous excitation by noise. In models for neural systems driven by noise,[11–13] networks with nonlocal connections between their elements exhibit a tendency towards relatively ordered oscillations in the population activity. The spatial structure of such noise-induced waves becomes less coherent with an increase in the fraction of nonlocal connections.[14] The combination of local connectivity with a small number of nonlocal connections allows a time-periodic localized external input to entrain the whole system much faster than in a purely local network; at the same time, the oscillations are much more coherent than in a truly random network.[15,16] Oscillatory activity at the population level has also been observed in small-world networks of binary McCulloch-Pitts neurons connected through both excitatory and inhibitory synapses.[17]

The dynamics of networks of excitable elements depend very strongly on both the range of the coupling and on the length of the refractory period relative to the time scales associated with propagation. In Ref. 18, the authors consider both scale free and random networks of three-state excitable neurons, and find that short refractory periods enable the propagation of sustained activity through the activation of short loops. These short loops are frequently found in many types of complex and random networks. In the noiseless case, they can only be activated if the initial conditions capture the broken symmetry associated with a choice of direction for the propagation of activity.[18] A related model of three-state excitable elements on a small-world network has been proposed to study the propagation of infectious diseases. In this model, the introduction of nonlocal connections was found to induce a transition to a state with coherent population oscillations.[19]

Most studies of the dynamics of complex networks have assumed that the interaction between node elements is bidirectional. This is a very reasonable assumption in the context of epidemic models,[19] and it would also apply to ensembles of neurons connected via gap junctions.[12] In the absence of noise, activity that is initiated through the excitation of individual neurons in an otherwise quiescent state results in wave fronts that propagate symmetrically in both available directions; if both local and nonlocal connections are bidirectional, all wave fronts are generated as symmetric pairs that annihilate upon collision. In such a system, persistent activity can only arise from initial conditions that suitably break the symmetry between the two possible directions for the propagation of activity.

In the cortical neural ensembles that motivate this work, the coupling between neurons is predominantly not bidirectional; most connections involve chemical synapses that transmit information from the presynaptic axon to the postsynaptic dendrite. In this scenario, it is more appropriate to consider directed networks with unidirectional connections. In previous work,[20] we have investigated networks in which the local connections are bidirectional, based on the assumption that the probability for reciprocal axo-dendritic connections is quite high for neurons in close proximity, while the nonlocal connections that provide shortcuts are unidirectional. This network is a modification of the by now classical small-world network,[21] in that the added nonlocal connections are directed. We found that just a few unidirectional shortcuts suffice for sustaining persistent activity, even when activity arises from localized excitations that do not break the symmetry between the two possible directions of propagation. However, as the density of shortcuts is increased, an increasing number of network configurations support only a brief burst of population activity after which the activity dies out. When the speed associated with the propagation of activity is low, this failure of network activity was found to be delayed and to occur only after many cycles of chaotic population bursts. The simplicity of this model allowed for an analytic description of the failure transition and for detailed numerical analysis.

The properties of this simple model[20] provide important insight into the phenomena found in simulations of more elaborate models motivated by specific biological systems. In Ref. 22, the connection between the topology of a neural network and its tendency towards epileptic seizures has been studied and related to the degree of recurrent connectivity in

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

different parts of the hippocampus. The origin of bursting behavior was addressed in Ref. 23. The common view on bursting behavior is that it arises when fast spiking drives a slow process, typically associated with slow kinetics, that in turn can shut off the spiking activity. However, no such slow kinetics are needed if the network displays small-world connectivity.[23] Both the seizing activity described in Ref. 22 and the bursting activity described in Ref. 23 find a common interpretation in the mechanisms that underlie the failure transition found in the simple model analyzed in our earlier work.[20] A rapid spread of activity followed by persistent oscillations has also been observed in recent chemical experiments based on the Belousov-Zhabotinsky reaction; in this system, unidirectional shortcuts were implemented through the photosensitive properties of the reaction.[24,25]

Here we build on our previous results, and present a detailed characterization of the persistent states and the dependence of their properties on the density of shortcuts; we also investigate the long chaotic transients and provide an explanation for the stretched exponential that characterizes their eventual failure. In Sec. II we define the model: a network of excitable integrate-and-fire neurons coupled via excitatory pulses in a small-world topology with unidirectional shortcuts. In Sec. III we discuss the persistent states and the crossover from persistent activity to failure for the case of rapidly propagating waves. In Sec. IV we analyze the exceedingly long chaotic transients in the regime of slowly propagating waves. In the concluding Sec. V we discuss our results in light of other work on neural networks with small-world topology.[4,22,23]

II. NEURAL MODEL AND NETWORK CONNECTIVITY

We consider a one-dimensional network: a ring of N identical integrate-and-fire neurons. The state of the ith neuron is described by its membrane potential V_i, which is dynamically controlled by the spiking activity of the neurons that project onto it,

$$\tau \frac{dV_i}{dt} = -V_i + RI_{\text{ext}} + g_{\text{syn}} \sum_{j=1}^{N} w_{ij} \delta(t - t_j - \tau_D). \tag{1}$$

In order to fully specify the dynamics of integrate-and-fire neurons, this set of N equations needs to be supplemented with a condition for spiking whenever the membrane potential reaches a specified threshold, and with a reset condition for the membrane potential after the emission of a spike,

$$V(t^+) = V_{\text{res}} \quad \text{whenever } V(t^-) = V_{\text{th}}. \tag{2}$$

In Eq. (1), τ is the membrane time constant, g_{syn} is the synaptic strength measuring the change in membrane potential due to each incoming spike, $w_{ij}=1$ or 0 indicates the presence or absence of a synaptic connection from neuron j to neuron i, t_j is the time at which neuron j fires a spike, I_{ext} is an external current, and R is the membrane resistance. The effective delay τ_D in the neuronal interaction includes both the time for the spike or action potential to propagate along the axon and the time needed for initiating the triggered action potential. When this latter time dominates over the axonal delay, the dependence of τ_D on the physical distance between the presynaptic and the postsynaptic neurons can be neglected. Postsynaptic currents due to synaptic activation are considered instantaneous and are therefore modeled as delta functions. The process of spike emission is described in Eq. (2): whenever the membrane potential of a neuron reaches the threshold value V_{th}, a spike is emitted and the membrane potential is reset to the value V_{res}. Without loss of generality, we write these two equations in terms of dimensionless quantities by measuring V_i and g_{syn} with respect to V_{res}, setting $V_{\text{res}}=0$ and $V_{\text{th}}=1$, and rescaling time by τ. (In Ref. 20, we used $\tau=10$.) Also, we replace RI_{ext} by the steady-state voltage V_∞; this is the asymptotic value that V_i would reach in the absence of synaptic input. We restrict the model to the case of excitable rather than spontaneously oscillating neurons by setting $V_\infty < V_{\text{th}}=1$, and consider only initial conditions such that at most a few neurons are triggered to spike while the rest of the network is in a quiescent state. In the absence of noise, this initial condition implies that neurons can only fire at times that are integer multiples of τ_D. Since the dynamical evolution of the membrane potentials can be integrated exactly between subsequent spikes, the time step for the numerical computations is taken to be $\Delta t = \tau_D$.

Cortical neurons often receive not only local input from nearby neurons but also input from some distant neurons through long-range projections. We mimic this heterogeneous connectivity through an extremely simplified network architecture: each neuron is bidirectionally connected to its $2k$ nearest neighbors, i.e., $w_{ij}=1$ for $|i-j| \leq k$, $j \neq i$, and unidirectionally connected to pN randomly chosen neurons. The parameter p thus indicates the density of additional unidirectional connections, as a fraction of the total number N of neurons.

The dynamics that arise from Eqs. (1) and (2) in the case of purely local connectivity ($p=0$) and in the absence of noise depend only on the interplay between the strength g_{syn} of the synapses, the number $2k$ of local connections per neuron, and the delay τ_D. The dependence on g_{syn} and τ_D is most easily illustrated for first-nearest-neighbor connectivity, $k=1$. In this case, if the presynaptic input is weak enough to satisfy $g_{\text{syn}} + V_\infty < V_{\text{th}}=1$, this presynaptic input is insufficient to cause a spike and the activity is not propagated. In contrast, if the presynaptic input is strong enough to satisfy $g_{\text{syn}} + V_\infty > V_{\text{th}}=1$, it results in a propagating wave of speed $1/\tau_D$. After spike emission, the voltage is reset to $V_{\text{res}}=0$, and the neuron is only ready to fire again after it has recovered to the extent that an input of magnitude g_{syn} is sufficient to trigger another spike. This recovery time is given by

$$T_R = \ln\left(\frac{V_\infty}{V_\infty + g_{\text{syn}} - 1}\right). \tag{3}$$

Note that T_R is not the intrinsic refractory period of the neuron, since this integrate-and-fire neuron can fire at arbitrarily large frequencies for sufficiently strong input g_{syn}.

Due to the bidirectionality of the local connections, the firing of each neuron not only triggers a spike in the neuron ahead of it in the direction of wave propagation, but also gives an input to the neuron behind it, which thus receives an input at a time $2\tau_D$ after its own firing. If $T_R > 2\tau_D$, this input is not sufficient to trigger a new spike, and the activity

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

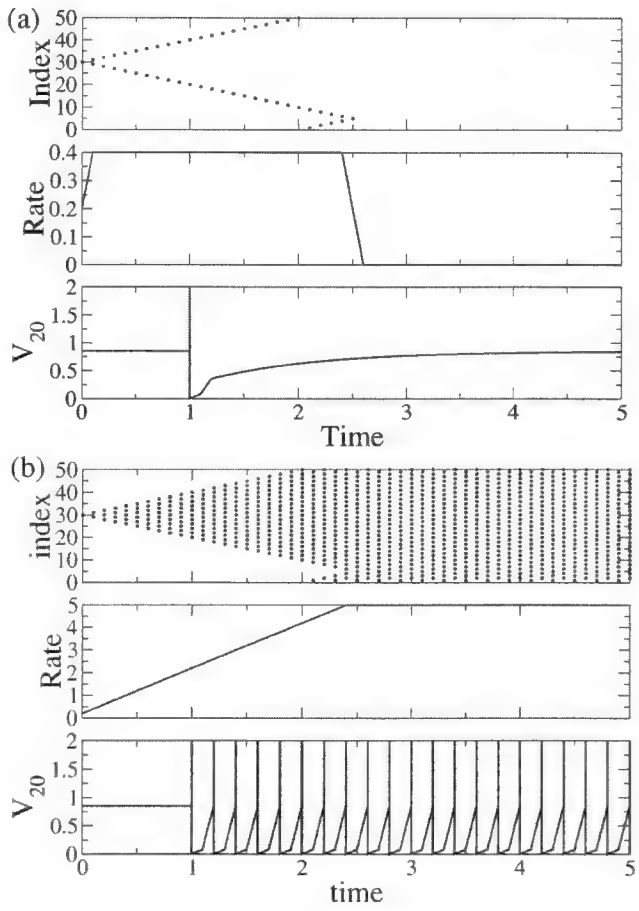

FIG. 1. Dynamics in a regular network of 50 neurons with first-nearest-neighbor coupling. Parameters are $V_\infty=0.85$ and $\tau_D=0.1$. (a) A weak synaptic coupling $g_{syn}=0.2$ results in a propagating wave. The two wave fronts meet and annihilate at time 2.5. Top: Raster plot showing spike times for each neuron. Middle: Population firing rate. Bottom: Membrane potential of neuron number 20. (b) A strong synaptic coupling $g_{syn}=1.0$ results in a wave that entrains all neurons in its path. The final state of the network consists of two synchronous groups of neurons firing out of phase with each other. Top, middle, and bottom panels as in (a). Note that the maximum possible firing rate is $1/\tau_D=10$.

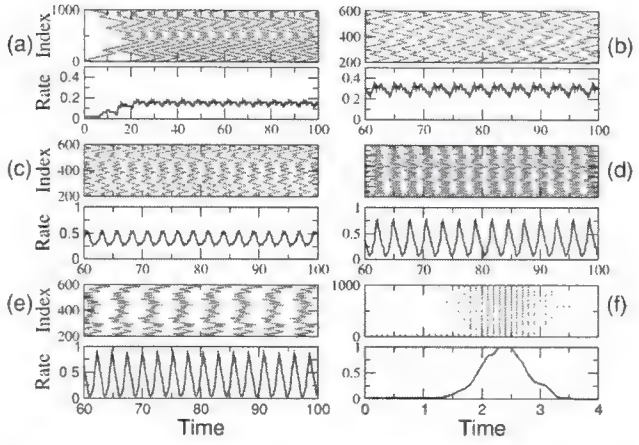

FIG. 2. Examples of network dynamics for $\tau_D=0.1$ and for different values of the density p of shortcuts. The values $p=0.01, 0.05, 0.10, 0.15, 0.20, 0.25$ correspond to panels (a)–(f), respectively. Each panel shows both spiking activity and population firing rate. Note different temporal scale on panel (f).

propagates away from the site of initiation as a wave to which each neuron contributes exactly one spike, see Fig. 1(a). However, if $T_R \leqslant 2\tau_D$, the wave front entrains all the neurons in its wake, eventually leading to synchronized activity of the whole network. In the absence of autapses, i.e., for $w_{ii}=0$, the network breaks up into two synchronous groups of neurons that fire out of phase with one another, see Fig. 1(b). In the general case, a neuron receives inputs from k neighbors as the wave approaches; the input coming from a neuron at a distance n is discounted by the factor $e^{-(n-1)\tau_D}$. A propagating wave can thus be sustained if $g_{syn}e^{\tau_D}\sum_{n=1}^{k}e^{-n\tau_D}+V_\infty > V_{th}=1$.

For simplicity, we will focus on the case of first-nearest-neighbor coupling in the regime in which waves of excitation propagate but do not entrain activity in their wake. This choice implies $k=1$ and constrains the allowable values of V_∞ and g_{syn}; we use $V_\infty=0.85$ and $g_{syn}=0.2$ unless otherwise noted. In this regime, the collision of two waves leads to their mutual annihilation and, after having fired in a propagating wave, a neuron can be triggered to fire by a single input of size g_{syn} after a time

$$T_R^{(1)} = \ln\left(\frac{V_\infty - g_{syn}e^{2\tau_D}}{V_\infty + g_{syn} - 1}\right). \tag{4}$$

This calculation includes the input received from the neuron ahead in the wave, at a time $2\tau_D$ after spiking.

The model specified by these assumptions together with Eqs. (1) and (2) is used here as a simple model for the generation and propagation of waves of activity in cortical tissue. As discussed in the following sections, the incorporation of random connections qualitatively alters the dynamics of the network and sustains a rich variety of spatiotemporal patterns.

III. THE ORDERED REGIME: ATTRACTORS AND FAILURE

The dynamics of the model depend on several parameters. Once the input current $V_\infty=0.85$ and the synaptic strength $g_{syn}=0.2$ have been fixed, the dynamics arising from Eqs. (1) and (2) are determined by the remaining two parameters: the fraction p of randomly placed shortcuts and the delay τ_D associated with the neural interaction.

The dynamics for $p\neq 0$ differ qualitatively from those for $p=0$. The presence of shortcuts allows the waves of excitation to be reinjected into portions of the network which have been previously excited. This process of reinjection may lead to persistent network activity, as shown in Fig. 2 for several values of the density p. As the waves spread outward from the initial site of activation, they encounter shortcut connections that inject activity elsewhere in the network. As for the $p=0$ case shown in Fig. 1(a), wave fronts that meet annihilate. After some time, the activity settles into a stable pattern in which the rates of wave generation and annihilation are balanced.

When averaged over time and across network configurations at a fixed value of the shortcut density p, the firing rate of these persistent states increases rapidly with p, and saturates around $p\sim 0.1$ [Fig. 3(a)]. This saturation is a conse-

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

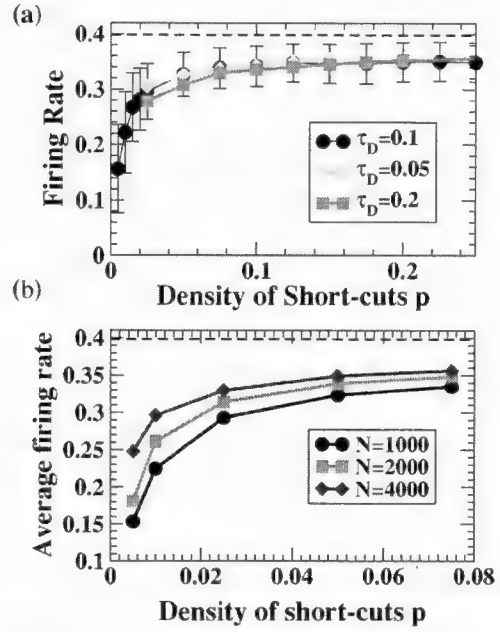

FIG. 3. Population firing rate, averaged over 1000 configurations for $p<0.2$ and over 5000 configurations for $p>0.2$. The maximal value of the firing rate, $1/T_R^{(1)}$, is indicated in both panels by a dashed horizontal line. (a) Firing rate as a function of the density of shortcuts p for different values of the delay τ_D for a network of size $N=1000$. Bars give the standard deviation of the firing rate across configurations for $\tau_D=0.1$. (b) Firing rate as a function of the density of shortcuts p for different values of the system size N with $\tau_D=0.1$.

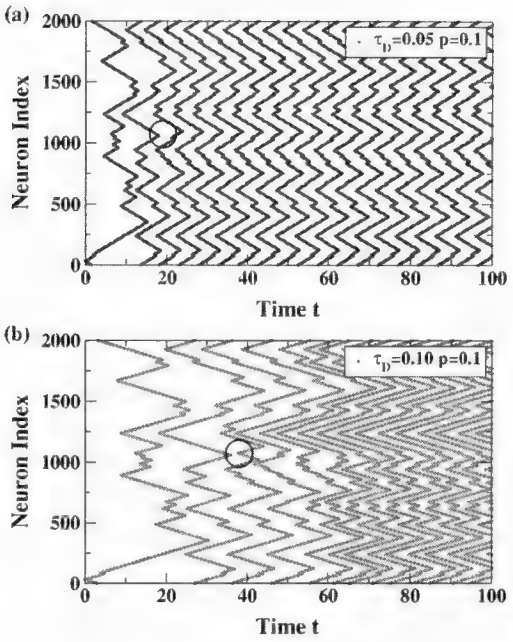

FIG. 4. Dynamics in a network of $N=2000$ neurons with a density of shortcuts $p=0.1$. The same network configuration and the same initial activation is used in both panels. The only difference is in the delay time: $\tau_D=0.05$ for the top panel and $\tau_D=0.10$ for the bottom panel. Note that more waves travel through the system when the delay is longer. Circles indicate a new wave spawned for $\tau_D=0.10$ but not for $\tau_D=0.05$.

quence of the neuron's finite recovery period $T_R^{(1)}$ of Eq. (4); its inverse is the maximal firing rate, indicated as a dashed line in both panels of Fig. 3. As shown in this figure, the firing rate comes very close to this maximal value. As the system size is increased, saturation is reached at smaller values of p [Fig. 3(b)].

The firing rate is essentially the inverse of the time between successive waves passing through a specific node in the ring of neurons. Thus, one may expect that decreasing the wave speed by increasing τ_D would reduce the mean firing rate. However, this is not the case. As shown in Fig. 3(a), the firing rate is quite insensitive to wave speed. The reason for this effect is illustrated in Fig. 4. The two panels show the spatiotemporal pattern of activity for the same network configuration and the same initial activation, but for different delay times τ_D. At the larger delay time, additional waves are excited through the shortcuts. An example of the appearance of such a new wave can be seen at $t=380\tau_D$, a time indicated by a circle in both panels of Fig. 4. At this time, a new wave is spawned by neuron 1071 for $\tau_D=0.1$ but not for $\tau_D=0.05$. These additional waves increase the firing rate in a manner that may even overcompensate for the reduced wave speed, as demonstrated by the higher density of waves for $\tau_D=0.1$ at times beyond $t=800\tau_D$.

Only a few shortcuts are present for low p, and many pathways leading to persistent activity consist only of large closed loops, which result in low firing rates. As p increases, the typical loop size decreases, and network configurations with only large loops and correspondingly low firing rates become increasingly unlikely. This effect is illustrated in the probability density function for the firing rate, shown in Fig. 5. As p increases, the distribution is shifted towards larger firing rates and it narrows substantially, reflecting a saturation close to the maximal firing rate set by the recovery time $T_R^{(1)}$. This maximal firing rate is marked by a dashed vertical line in Fig. 5.

For small values of p, the state of persistent activity settles into a periodic pattern whose oscillations increase in amplitude with increasing p. This effect is clearly illustrated in panels (a)–(e) of Fig. 2. A quantitative description of this effect is shown in Fig. 6, in which the standard deviation of the firing rate, averaged over a large number of network configurations that exhibit persistent activity for a given value of p, is used to characterize the amplitude of the oscillations. An additional feature of this regime, also apparent from panels (a)–(e) of Fig. 2, is that the time for these oscil-

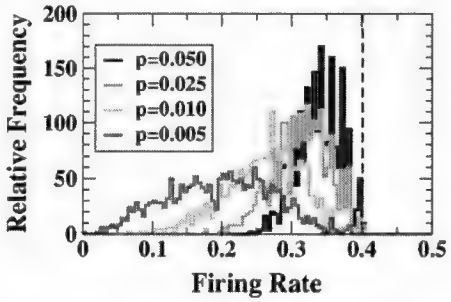

FIG. 5. (Color) Probability distribution of firing rates for several values of p, for networks of $N=2000$ neurons with $\tau_D=0.1$. As p increases, the distributions narrow in width as they shift towards the maximal firing rate $1/T_R^{(1)}$, indicated by a dashed vertical line.

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

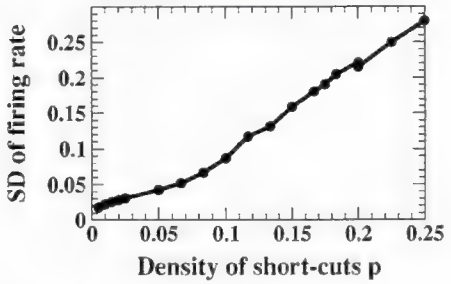

FIG. 6. Amplitude of oscillatory activity, measured through the standard deviation of the firing rate, as an increasing function of the shortcut density p. The standard deviation is obtained as an average over 5000 network configurations that exhibit persistent activity for $\tau_D=0.1$; $N=1000$.

lations to become established after the excitation of a single neuron decreases with increasing p. This property is in agreement with the results of Refs. 15 and 16, which found that, when excited through a small cluster of driven oscillatory neurons, small-world networks of either Hodgkin-Huxley or FitzHugh-Nagumo neurons are entrained much more quickly than regular networks (with no shortcuts) of the same types of neurons.

For larger values of p, the activity patterns can be quite complicated. In this regime, as in the small p regime, all neurons get excited during an oscillation cycle of network activity. However, in this regime not all neurons and not all connections between them are necessary for the persistence of activity, as illustrated in Fig. 7. The full raster plot of network activity shown in the top panel depicts the spikes from all neurons (black dots). This plot also identifies the spikes from those neurons that are essential for sustainability (red dots). These neurons provide a pathway for recurrence,[26] and they are found as follows. At an arbitrary time once the steady state has been reached, all neurons that fire at that time step are labeled. A backwards search is then performed for preceding ancestors of these labeled neurons,

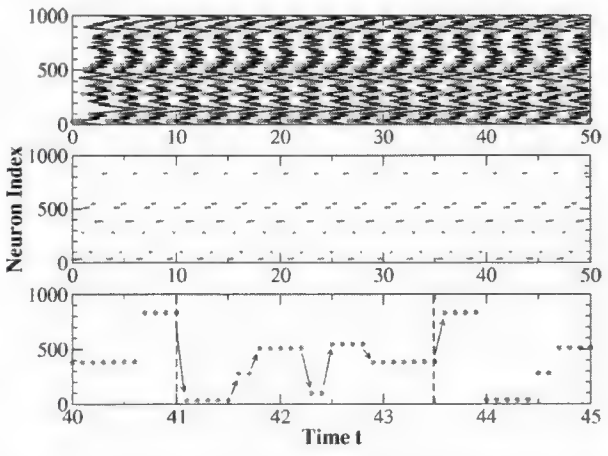

FIG. 7. Spiking activity in a network of $N=1000$ neurons with a density of shortcuts $p=0.1$ is shown in the top panel for $\tau_D=0.1$. The activity of neurons that are part of the backbone pathway is indicated in red. The middle panel shows the spiking activity within the backbone. The enlargement shown in the bottom panel displays periodic activity, with a period longer than the recovery time $T_R^{(1)}=2.494$ (delimited by dashed vertical lines). See text for details.

i.e., those presynaptic neurons that fired one delay τ_D ago and triggered the activity of the labeled neurons. These one-step ancestors are labeled in turn, and the process is iterated backwards in time until $t=0$ is reached. The pattern of labeled neurons quickly converges to a small subset, as shown in the middle panel of Fig. 7. Only the neurons in this reduced subset contribute to the persistence of the pattern; the remaining neurons could be cut out of the network without destroying the persistent activity. The pattern of activity along this backbone, shown in the bottom panel of Fig. 7, is periodic in time. The period is only slightly longer than the recovery time $T_R^{(1)}$, indicated by the dashed vertical lines (note the expanded time scale in this panel). Such backbone pathways have also been identified in an experimental study of the excitable Belousov-Zhabotinsky reaction;[25] in this photosensitive system, unidirectional shortcuts were implemented through local optical excitation.

A striking feature of this regime is that a network configuration capable of sustaining persistent activity displays an extraordinarily large number of different attractors. To assess the number of coexisting attractors we focus on the regime of low shortcut density, where all solutions are periodic. Each solution can be characterized by its period, its mean firing rate, and the standard deviation of the firing rate. A labeling of attractors based only on these three measures is likely to underestimate the total number of attractors; however, it suffices to find a large number of them, as shown in Fig. 8. This figure illustrates attractor multiplicity for a randomly chosen network configuration of $N=1000$ neurons with $p=0.05$ and $\tau_D=0.1$. Each of the 471 distinct stable patterns of persistent activity identified in Fig. 8(a) arises from an initial condition in which only one neuron is activated. As shown in Fig. 8(b), which shows the number of initial conditions that lead to a solution with a given period, many of these stable patterns have the same period (note the logarithmic scale). In contrast, the standard deviation of the firing rate [Fig. 8(c)] shows great variability, reflecting different temporal evolutions of the firing rate within a period.

A large number of attractors is typical for these networks. Figure 9 shows the number of attractors as a function of network size; black circles represent data points obtained as averages over 20 different network configurations of a fixed size. Most attractors have quite small basins of attractions: within the restricted set of initial conditions in which only a single neuron is excited, most attractors can be reached from only one such initial condition (red squares). The overall number of attractors increases roughly linearly with system size for large values of N, as does the number of attractors with different periods (blue diamonds), but the latter is about an order of magnitude smaller than the total number of attractors. The identification of attractors requires that networks achieve a steady state of persistent activity, but since the duration of the transients grows with system size, the computation time grows faster than N^2. This precludes us from simulating significantly larger system sizes than shown in Fig. 9; such data would be necessary to reliably estimate scaling relations between network size and number of attractors. (For $N=4000$ the computation takes over 2 weeks on a desktop PC.) Preliminary computations with more general

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

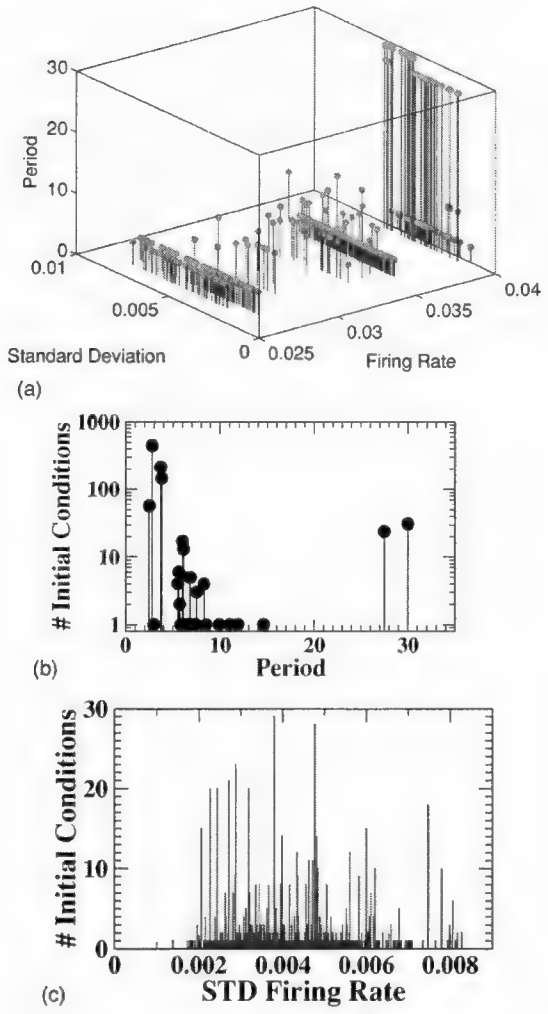

FIG. 8. Large number of attractors for a specific network configuration, with $N=1000$ neurons, $p=0.05$, and $\tau_D=0.1$. (a) The mean firing rate, its standard deviation, and the period of all 471 distinct attractors evoked by initial conditions in which only one neuron is activated. (b) Number of initial conditions that lead to an attractor with a given period; many attractors, although distinct, have the same period (note the logarithmic vertical scale). (c) Number of initial conditions that lead to an attractor with a given standard deviation of the firing rate.

initial conditions reveal many more attractors than those shown in Fig. 8. Thus, while the restricted initial conditions that give rise to Fig. 9 suggest an almost linear increase in the number of attractors with system size, the full number of attractors may grow substantially faster.

At this point, the origin for this exceedingly large number of different attractors is not clear. In all-to-all coupled oscillator systems, factorially large numbers of attractors are due to the permutation symmetry associated with the global coupling.[27] The small-world networks investigated here do not possess such symmetry. In this case, the large number of attractors found for a specific network configuration is likely to be due to a combination of geometrical and dynamical causes: the coexistence of many backbones that can independently support sustained activity, and the variability in membrane potentials that can result in distinct but not too different firing patterns. While we find transitions between different attractors in the presence of finite-amplitude

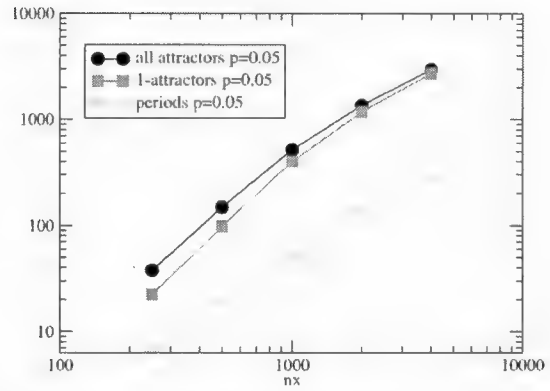

FIG. 9. Number of attractors: all attractors (circles), attractors with basin of attraction of size one (squares), and attractors with different periods (diamonds; note different vertical scale on the right) as a function of network size N.

noise,[26] we have not investigated whether this system exhibits the extreme noise sensitivity found in the case of all-to-all coupled oscillators, where it is due to the crowding of large numbers of attractors.[27]

As the density p of shortcuts increases, the distance traveled by the waves before encountering the entrance to a shortcut decreases. Consequently, waves of excitation spread throughout the network more rapidly, as shown in the progression of spatiotemporal patterns of Figs. 2(a)–2(e) for increasing values of p. If p becomes too large, activity spreads too fast and it quickly dies away, as shown in Fig. 2(f). Since the shortcuts are randomly placed, different network configurations will exhibit different dynamics. Thus, while the overall likelihood of persistent activity decreases with increasing p, the actual network dynamics depend on the particular network configuration.

The mechanism that leads to the extinction of network activity is easily elucidated.[20] Once a neuron has emitted a spike, its membrane potential is reset to a fixed value, chosen here to be $V_{\mathrm{res}}=0$. While the membrane potential of the neuron recovers towards its resting value V_∞, activity spreads through the network, eventually finding its way back. Once this occurs, the neuron receives synaptic input equal to g_{syn}. This input will be sufficient to trigger a spike only if this neuron has recovered sufficiently. It is therefore clear that as the number of shortcuts is increased and activity spreads more rapidly, the network is less likely to sustain persistent activity. Whether or not this mechanism of premature return will lead to the extinction of activity in a given network depends on its particular configuration. Many different network configurations for a given value of p have been simulated in order to measure the fraction that fail to sustain persistent activity. This probability of failure is shown as a function of p for different values of the system size N in the left inset of Fig. 10. In agreement with our intuitive argument, the probability that a network drawn at random fails to sustain activity increases with increasing p. In fact, there is a sharp crossover from a low p regime characterized by the ability to sustain persistent activity into a large p regime in which activity will always fail. The transition between these

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

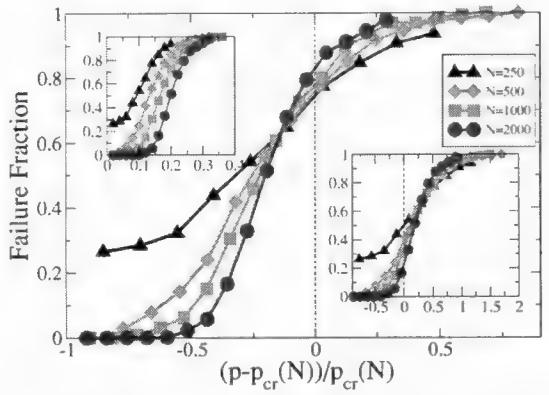

FIG. 10. Probability of failure as a function of shortcut density for different values of the system size N. The density of shortcuts is scaled by p_{cr} from Eq. (8). Left inset: unscaled data. Right inset: data scaled by p_{cr} from Eq. (7).

two regimes occurs at higher values of p as the system size N increases.

As discussed in Ref. 20 this transition can be captured in a mean-field approximation in which the return time is assumed to be identical for all neurons. In this approximation, the maximum return time T_A needed for the activity to traverse the entire network can be expressed as a function of only p and N. Setting this time to be equal to the recovery time yields an upper bound for the critical density of shortcuts at which a transition from persistent activity to failure occurs,

$$T_A(p_{cr}) = T_R^{(1)}. \tag{5}$$

An approximate form for T_A is easily derived. Assume an initial condition in which a single neuron fires at time $t=0$. Given a density p of shortcuts, the entrance to a shortcut will be typically encountered after $1/p$ neurons have fired, which occurs after a time $\tau_D/2p$; the factor of 2 is due to the two wave fronts that emerge from the initially activated neuron and propagate in opposite directions. Due to the activity injected through the shortcut, four wave fronts are now propagating through the system; $2/p$ neurons will fire during the subsequent time interval of duration $\tau_D/2p$, at the end of which a new shortcut entrance will typically be found and two more wave fronts will be generated. The process is iterated, with $2^{k-1}/p$ neurons firing during the kth cycle. It takes n cycles to ensure that all neurons have fired, with n such that

$$\sum_{k=0}^{n-1} 2^k = pN, \tag{6}$$

which leads to a total time

$$T_A(p) \equiv n\frac{\tau_D}{2p} = \tau_D \frac{\ln(1+pN)}{2p \ln 2}. \tag{7}$$

This is a purely geometric result for the time it takes for activity to traverse the entire extent of the network; this time is related trivially to the largest distance in the network. The geometric mean-field properties of small-world networks have been analyzed by Newmann, Moore, and Watts,[28] who used a continuum limit to calculate the fraction of a small-world network that is covered by starting at a single point and extending outwards a distance r in both directions; this is equivalent to following the spread of waves of activation during a time $r\tau_D$ for the neural network considered here. The calculation in Ref. 28 takes into account two effects that were omitted from their earlier calculation[29] and from our derivation, Eq. (7). First, as the network is covered, a shortcut might lead to a part of the network that has already been traced over; this contribution should not be counted. In the neural scenario, this is equivalent to an attempt at injecting activity into a neuron that has already fired but not yet recovered to the point where it can fire again. Such a shortcut does not contribute to sustained activity. Second, when two covering fronts meet, they stop and no longer contribute. In the neural network, this corresponds to activity wave fronts that meet and annihilate. The incorporation of these two additional mechanisms leads to a two-component model that correctly describes both the covered fraction of the network and the number of fronts.[28] The result, when applied to the neural network of Eq. (1), yields

$$\sqrt{\left(1+\frac{4}{pN}\right)} \tanh\left[\sqrt{\left(1+\frac{4}{pN}\right)}\frac{pT_A(p)}{2\tau_D}\right] = 1. \tag{8}$$

The identification $T_A(p_{cr}) = T_R^{(1)}$ then yields the mean-field estimate $p_{cr}^{(MFT)}$ for the density at the failure transition.

The main panel in Fig. 10 shows the failure rates as a function of shortcut density, with p rescaled by the critical density $p_{cr}^{(MFT)}$ that follows from Eq. (8). All rescaled curves intersect at a common value of p, which defines the transition point at p_{cr}. The mean-field theory yields an upper bound, $p_{cr} < p_{cr}^{(MFT)}$. It is interesting to note that a similar rescaling using the estimate for p_{cr} that results from Eq. (7) also produces a family of rescaled curves that intersect at a common value of p, as shown in the right inset of Fig. 10. While $p_{cr}^{(MFT)}$ overestimates the true p_{cr} at which the rescaled curves cross, the estimate based on Eq. (7) underestimates the true p_{cr}. In an earlier report[20] we incorrectly used T_R instead of $T_R^{(1)}$ as the recovery time in Eq. (8). The corresponding rescaled curves intersected at zero. We cannot tell whether this agreement with the true p_{cr} was just a coincidence or the consequence of a subtle cancellation between factors leading to overestimation or underestimation. In any case, it is worth emphasizing that all three methods, although only approximate, succeed in rescaling the curves in the left inset of Fig. 10 so as to obtain a unique crossing. This is due to the fact that these phenomenological approaches result in estimations for the critical density p_{cr} with very similar dependencies on system size N; these estimates seem to provide a good approximation to the true dependence on system size.

The failure transition occurs at the value of p for which the geometrical quantity $T_A(p)$ equals $T_R^{(1)}$. This recovery time is thus the central quantity that determines persistence or failure for a given shortcut configuration. It is worth pointing out that this recovery time is not the same as the absolute refractory period T_r. During the refractory period, neurons are inhibited from receiving synaptic input, while the mem-

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

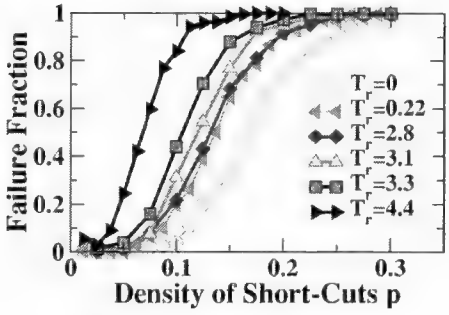

FIG. 11. Dependence of failure probability on shortcut density for different values of the refractory period T_r. Shown are averages over 500 network configurations, with $N=1000$, $\tau_D=0.1$, and $T_R=2.83$.

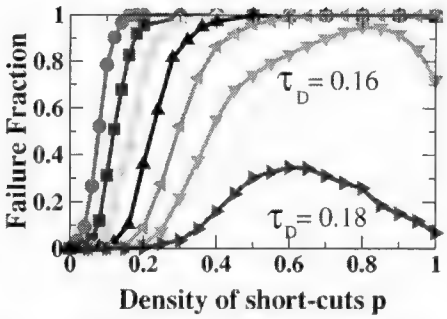

FIG. 12. Dependence of failure probability on shortcut density for different values of the delay $\tau_D=0.06,0.08,0.10,0.12,0.14,0.16$. Shown are averages over 2000 network configurations for $N=1000$. Note the nonmonotonic character of the curves for large enough delay. In the large p regime, failure probability is estimated within a finite time $T^*=100$.

brane potential still relaxes towards its resting value. While T_r is an intrinsic property of individual neurons, $T_R^{(1)}$ depends strongly on the strength of the synaptic coupling. The curves in Fig. 11 show the fraction of network configurations that fail to sustain persistent activity for different values of the refractory period T_r. For $T_r<2\tau_D$, the input that neurons in the wake of a propagating wave receive at a time $2\tau_D$ after their own firing is unaffected, and the relevant recovery time is $T_R^{(1)}$ [cf. Eq. (4)]. The corresponding curves are as those for $T_r=0$, shown in Fig. 10. The input received at $2\tau_D$ after firing is suppressed for $T_r>2\tau_D$. In this low p regime, the persistent states do not depend on neurons receiving additional inputs before the one that triggers a spike, and the failure transition is controlled by T_R [cf. Eq. (3)]. The failure transition is thus independent of T_r as long as $2\tau_D<T_r<T_R$. It is only for $T_r>T_R$ that the refractory period suppresses relevant input to the neurons and affects the failure transition. This is illustrated in Fig. 11, where the refractory period T_r is seen to have little effect on persistent activity when it takes values between $2\tau_D=0.20$ and $T_R=2.83$.

In our analysis of persistent activity and transition to failure, we have also considered the possibility of an upper bound in the length of allowed shortcuts. We summarize our results without showing the corresponding numerical data: the behavior of the system is qualitatively unchanged as long as this upper bound exceeds a threshold value; below threshold, the network is essentially only locally coupled and the failure probability rapidly approaches one. Other modifications to the distribution of shortcut lengths might include allowing for a nonuniform distribution. It is unclear how a nonuniform distribution of shortcut lengths would affect the dynamics. However, it has been shown that one-dimensional networks maintain a small-world structure if the distribution of shortcut lengths is power-law with power <2.[30] We assume that the regime we have discussed in this section will also be present in such a scenario.

IV. THE DISORDERED REGIME: CHAOTIC TRANSIENTS IN SLOW WAVES

For small values of the density of shortcuts p, a small-world network of integrate-and-fire neurons is quite likely to sustain persistent activity. As discussed in the preceding section, the spatiotemporal pattern of network activity in this regime is highly regular and most often periodic, despite the complex, heterogeneous topology of the network itself. As p increases, a transition takes place: more and more network configurations exhibit activity which peaks and then shuts down. Interestingly, for large enough τ_D there is an additional change in network dynamics as p increases. As shown in Fig. 12 for $\tau_D=0.16$ and 0.18, the failure probability initially increases with p, as for low values of τ_D, but it then turns downwards again as p increases further. To understand this reentrant phenomenon, we first analyze the spatiotemporal dynamics characteristic of this seemingly persistent activity at large p.

Examples of network dynamics for slow waves ($\tau_D=0.16$) are shown in Fig. 13. For values of p below or near the theoretical transition to failure [Figs. 13(a)–13(c)], the activity is similar to that shown for fast waves ($\tau_D=0.10$) in Fig. 2. However, in the reentrant regime [Fig. 13(d)], the activity is chaotic [cf. Fig. 15(a) below] and the population firing rate exhibits irregular peaks that reflect near-synchronous activity involving a large fraction of the network.

A detailed, quantitative analysis for $\tau_D=0.18$ shows that the change in behavior occurs already before the maximum of the failure probability curve. For fast waves, corresponding to small τ_D, the amplitude of the oscillations in the population activity, as measured by the standard deviation of the firing rate, was found to increase monotonically with p (cf. Fig. 6). However, for slow waves, corresponding to large τ_D, this amplitude of oscillations is nonmonotonic and decreases over the range $0.4\leqslant p\leqslant 0.6$ (Fig. 14). It is instructive to compute the spectral entropy of the population firing rate,

$$\mathcal{S}=-\sum_{\omega} P(\omega)\ln P(\omega), \tag{9}$$

which measures the number of significant peaks in the power spectrum $P(\omega)$. The spectral entropy exhibits a significant increase over the same range $0.4\leqslant p\leqslant 0.6$ [Fig. 15(b)], indicating an increase in the complexity of the dynamics. The variability of the spectral entropy across network configurations with the same density of shortcuts exhibits a broad maximum in the same range of p values, and reaches very small nonzero values (about 0.06) in the strongly chaotic regime. The detailed evolution towards chaotic dynamics,

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

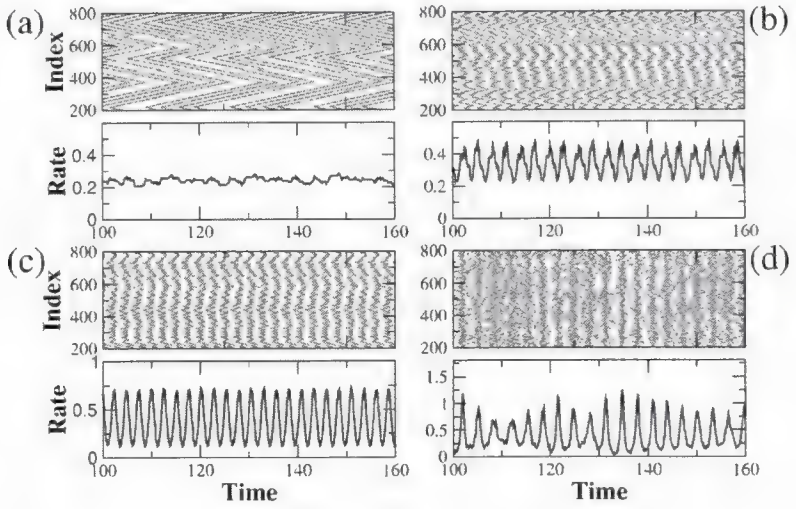

FIG. 13. Examples of network dynamics for $\tau_D=0.16$ and for different values of the density p of shortcuts. The values $p=0.01$, 0.2, 0.4, and 1.0 correspond to panels (a)–(d), respectively. Each panel shows both spiking activity and population firing rate. The reentrant activity in (d) is noisy and exhibits synchronized population bursts.

which depends on the specific configuration of shortcuts, is not investigated in further detail in this work.

What underlies the emergence of reentrant activity with increasing p? The answer lies in the interplay between network topology and the delay τ_D. The mean-field model that provides a description of the failure transition, Eq. (5), is based on the assumption that the maximal firing frequency of each neuron is limited by the recovery time $T_R^{(1)}$. However, a small-world network constructed by adding shortcuts allows for neurons to receive more than one incoming shortcut. The probability of such configurations is small at low p; in this regime, the fraction of neurons with two incoming shortcuts is given by $s_2 \sim p^2/2$,[26] which is indeed negligibly small for $p<1$. (We assume that the number of neurons with more than two shortcut inputs can be ignored and we compute the most likely rather than the expected value of the number of neurons with two shortcut inputs.) Mean-field results that ignore such configurations describe the failure transition provided it occurs at $p_{cr}<1$. However, as p approaches 1, i.e., for $|p-1| \ll 1$, the fraction of neurons with two incoming shortcuts becomes $s_2=1-\sqrt{2}/2+(p-1)/2$, which is nearly 0.3 for $p=1$. In this regime, a significant fraction of neurons is likely to receive several synaptic inputs during one cycle of network activity. Such neurons would not be constrained by the recovery time $T_R^{(1)}$, but would rather be primed to fire earlier, potentially allowing the activity to persist where it otherwise would fail. A neuron that has received n inputs at times t_r, $r=1,\ldots,n$, since its last firing has a recovery time $T_R^{(n)}$ given by

$$T_R^{(n)}(t_1,\ldots,t_n)=\ln\left(\frac{V_\infty-g_{\mathrm{syn}}\Sigma_{r=1}^{n}e^{t_r}}{V_\infty+g_{\mathrm{syn}}-1}\right). \tag{10}$$

Note that Eq. (10) reduces to Eq. (3) for $n=0$ and to Eq. (4) for $n=1$ (with $t_1=2\tau_D$). In general, $T_R^{(n)}$ depends on the specific firing times of the n neurons that provide inputs through shortcuts; these times depend in turn on the details of activity propagation in each specific network configuration. However, since integrate-and-fire neurons become increasingly sensitive to their input as time passes after their firing, the value of $T_R^{(n)}$ in Eq. (10) is bounded below by $T_{R,\min}^{(n)}$, which

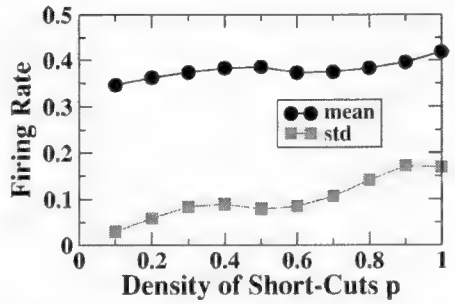

FIG. 14. Average population activity (mean firing rate) and amplitude of its oscillations (standard deviation of the firing rate) as a function of the shortcut density p. Note the nonmonotonic behavior of these curves. Data obtained as an average over those configurations (out of 200) for which the activity persists for at least 15 000 steps for $\tau_D=0.18$ and $N=1000$.

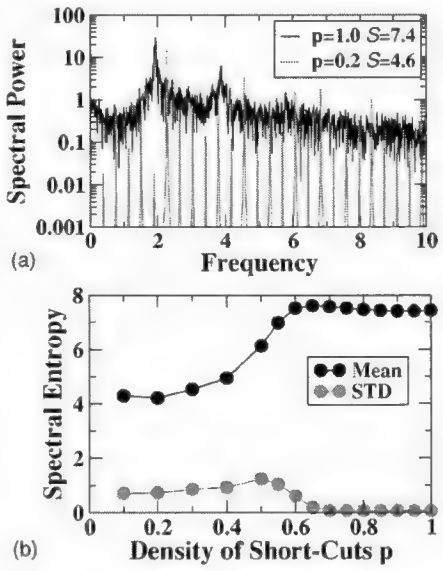

FIG. 15. Temporal complexity of activity patterns. (a) Representative power spectra for $p=0.2$ ($S=4.6$) and $p=1.0$ ($S=7.4$). (b) Spectral entropy S (mean and standard deviation across 200 configurations) for $\tau_D=0.18$ and $N=1000$, over 15 000 time steps.

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

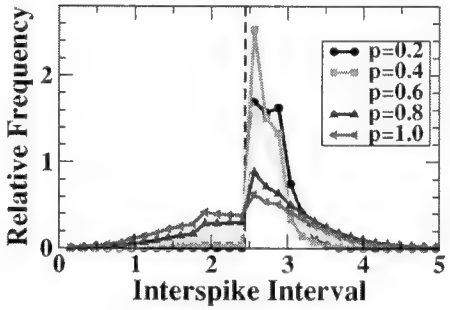

FIG. 16. Distribution of interspike intervals (ISI) for $\tau_D=0.16$ and different values of the shortcut density p. For low p, all allowed ISI are above the recovery period $T_R^{(1)}$ (dashed vertical line). For $p \geq 0.6$, multiple inputs result in ISI below $T_R^{(1)}$.

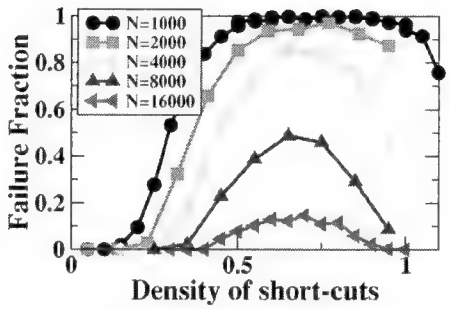

FIG. 17. Failure probability as a function of shortcut density p for $\tau_D=0.14$ and different values of system size N. Shown are averages over 400 configurations; failure probability is estimated within a finite time $T^*=28$. Note the prevalence of prolonged activity in larger systems.

occurs when all n inputs coincide at $T_{R,\min}^{(n)}$ itself. This lower bound is given by

$$T_{R,\min}^{(n)}(t_i = T_{R,\min}^{(n)}) = \ln\left(\frac{V_\infty}{V_\infty + n g_{\text{syn}} - 1}\right). \tag{11}$$

For small enough p, the interspike intervals (ISI) of almost all neurons in almost all network configurations are bounded below by $T_R^{(1)}$; this property has allowed us to calculate the time for activity to spread throughout the whole network using a purely geometric approach. For higher values of p, there may be a subset of neurons with shorter allowable ISI. However, many neurons will still receive only one input per cycle, and the frequency of their spiking activity should reflect this fact. The distribution function for the ISI shown in Fig. 16 supports this argument. Fast spiking activity with ISI$<T_R^{(1)}$ occurs appreciably only for $p \geq 0.6$, and becomes both more common and faster with increasing shortcut density. Our earlier analysis[20] showed that the spikes with ISI$>T_R^{(1)}$ occur in population bursts, with no such spikes in between bursts. In the absence of other spikes, the activity would die out during these intervals between bursts. However, the fast spiking neurons that receive multiple inputs via multiple incoming shortcuts are primed to carry over the activity during the intervals between bursts; their fast spiking sustains activity during the time needed for the slow spiking neurons to recover.

Long delays τ_D contribute in several ways to bridging the periods of low activity between bursts. For larger τ_D, the failure transition is shifted towards larger shortcut densities, thus significantly enhancing the number of neurons that receive multiple shortcut inputs. At the same time, in order to bridge the time between the return time T_A and the recovery time $T_R^{(1)}$ of the slow spiking neurons, fewer fast spiking neurons are needed if the delay is longer. Moreover, the exponential recovery of integrate-and-fire neurons towards their resting potential implies that later inputs have a stronger impact on the recovery period than earlier ones [cf. Eqs. (10) and (11)]. With increasing τ_D, all inputs are shifted to later times relative to the most recent spike of the postsynaptic neuron; this shift significantly reduces the recovery time. However, an increased delay τ_D is not necessary to establish a regime of prolonged activity. As shown in Fig. 17 for fixed delay $\tau_D=1.4$, increasing the system size from $N=1000$ to $N=16\,000$ shifts the failure transition to sufficiently large p that the number of neurons with multiple inputs is sufficient to bridge the gap between bursts of slow spiking activity, even for this shorter delay.

For low values of p, the spatiotemporal dynamics are most often periodic. In those cases the dynamics can truly be called persistent. For $p \gtrsim 1$, the chaotic nature of the dynamics precludes such a clear assessment; in fact, failure is possible even after very long times. In this regime, prolonged activity relies on bridging the quiescent period between bursts of slow spiking activity through fast spiking neurons that receive multiple shortcut inputs; it is necessary that these shortcuts are actually activated at suitable times during the quiescent part of the cycle. Thus, while in one cycle the activity during the burst may have excited such a pathway, the different activity pattern in the next burst may fail to do so; the activity could then die out. Indeed, we find that the prolonged activity characteristic of large shortcut densities eventually fails for essentially all configurations. Examples of such long-lived transients are shown in Fig. 18, where population firing rates are shown for a fixed network configuration of shortcuts at increasing values of τ_D.

The value of τ_D has a strong influence on the duration of the transient. As shown in Fig. 18, the overall trend is for the lifetime of the transient activity to increase with increasing delay (note the change in temporal scale from panel to panel). However, the actual dependence on τ_D is more subtle,

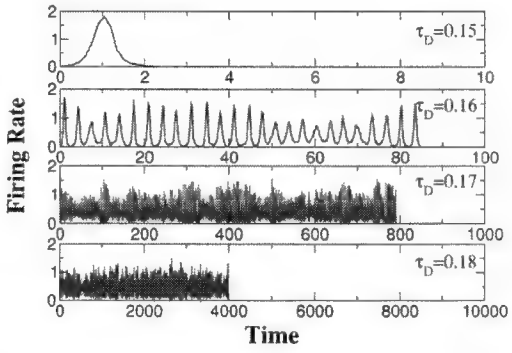

FIG. 18. Population firing rate for four values of the delay τ_D, with $p=1$. The same network configuration and the same initial activation is used in all four panels. Failure tends to occur later as τ_D increases; note the change in temporal scale from panel to panel.

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

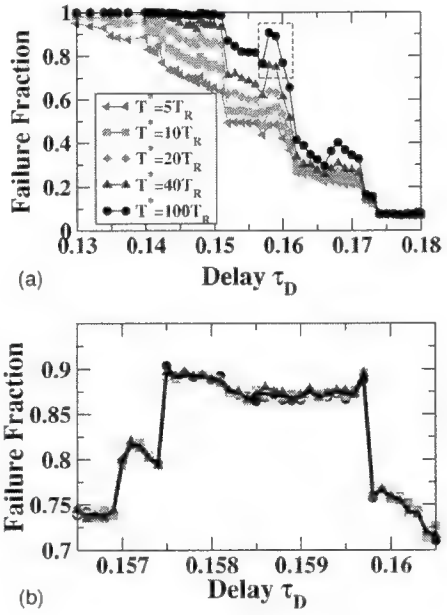

FIG. 19. (a) Failure rates for different values of the final time $T^*=5T_R$, $10T_R$, $20T_R$, $40T_R$, and $100T_R$. Shown are averages over 2000 configurations with $N=1000$ and $p=1$. (b) Enlarged view of the data within the box in (a) for $T^*=100T_R$ illustrates the fine structure in the failure rate. Shown are four runs based on averages over 8000 configurations each; the black line is the average over the four runs.

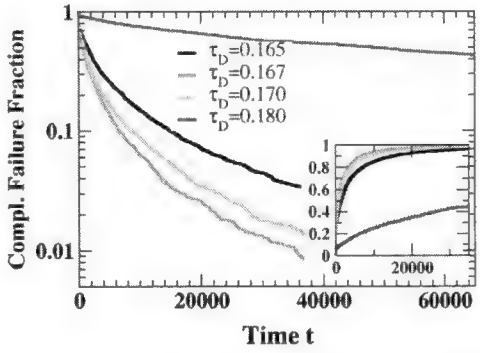

FIG. 21. (Color) Complementary failure fraction $1-\mathcal{F}$ as a function of time for different values of the delay τ_D. Note the logarithmic scale on the vertical axis. Inset: Failure fraction $\mathcal{F}$ for the same data; the vertical scale is now linear.

as shown in Fig. 19. The fraction of network configurations for which the prolonged activity fails before a specified time T^* is reached exhibits a surprising degree of structure in its dependence on τ_D. Most surprising is the finding that an increase in τ_D does not always decrease the probability of failure, but can in fact enhance it. These changes can occur over very small intervals in τ_D, as shown in the expanded window of Fig. 19(b). This fine structure is reminiscent of resonances, although these are more like antiresonances; the values of τ_D within these windows are in some sense optimal for escaping from prolonged activity. While details of the mechanism underlying this structure are not yet understood, it is clear that the dependence on τ_D reflects the significance of the ratio τ_D/T_R. This effect is illustrated in Fig. 20, which shows the dependence of the failure probability on τ_D for two network sizes: $N=500$ and $N=1000$. While the probability of failure is higher overall for the smaller network, the location of the "resonant" windows in τ_D does not show much dependence on system size. However, even a small reduction in recovery time from $T_R=2.83$ to $T_R=2.79$, due to an increase in synaptic strength from $g_{syn}=0.200$ to $g_{syn}=0.202$, results in a significant shift of the "resonant" windows towards lower values of τ_D.

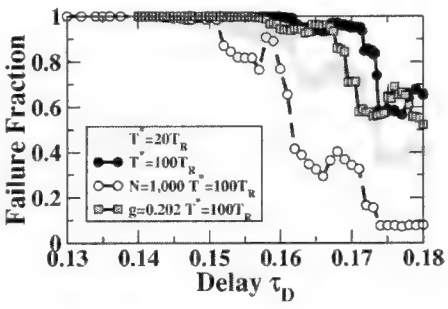

FIG. 20. Failure rates averaged over 8000 network configurations with $N=500$ and $p=1$. Data are shown for both $g_{syn}=0.200$ (circles) and $g_{syn}=0.202$ (squares). Data for $N=1000$ (cf. Fig. 19) shown for comparison.

To assess whether any significant fraction of the network configurations sustains truly persistent activity, we consider the failure fraction $\mathcal{F}$ as a function of the final time T^*, so as to extrapolate to $T^*\to\infty$. The complementary fraction $1-\mathcal{F}(T^*)$ of network configurations that sustain persistent activity up to a time T^* is shown for different values of the delay τ_D in Fig. 21. As anticipated from the nonmonotonicity in Fig. 19, the fraction of failing network configurations is largest for $\tau_D=0.167$; a very rapid drop in failure rate occurs from $\tau_D=0.17$ to $\tau_D=0.18$. When considered as a function of time for fixed τ_D, the behavior of the curves in Fig. 21 indicates that the decay is not exponential.

To obtain an approximate analytic form for the failure fraction $\mathcal{F}(T^*)$, let us consider a specific network configuration for a given value of p. The duration of the activity until failure will then depend on the initial condition. In this numerical experiment we choose a specific network configuration and consider 2000 different initial conditions; to reduce the computational effort these simulations are done for a smaller system with $N=200$. The initial conditions are random and given by $V_i=V_0+\xi_i$, $i=1\ldots N$, with $V_0=0.85$ and ξ_i drawn from a uniform distribution in the interval $[-0.5, +0.5]$. A failure time is determined for the activity triggered by each initial condition. The resulting distribution of failure times exhibits exponential behavior of the form $\beta e^{-t/T}$ for large times; this allows the extraction of a characteristic failure time T associated with this network configuration. An exponential distribution of failure times suggests that the chaotic dynamics effectively lead to a fixed probability for the activity to die out after each population burst.

The characteristic failure time T is then computed for many network configurations of the same size and same shortcut density, leading to a distribution $\rho(T)$ of characteristic failure times, as illustrated in Fig. 22. The large-T behavior of this distribution is well fit by an exponential decay,

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

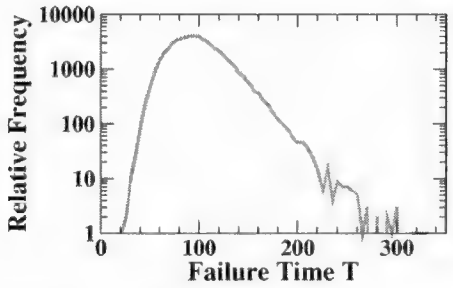

FIG. 22. Distribution of characteristic failure times T, from 50 000 network configurations with $N=200$, $p=1$, and $\tau_D=0.16$. Note the logarithmic scale on the vertical axis.

$$\lim_{T\to\infty} \rho(T) \propto e^{-\alpha T}. \tag{12}$$

The distribution $\rho(T)$ of characteristic failure times determines the average number of failures expected to have occurred by time T^*,

$$\mathcal{F}(T^*) = \int_0^\infty dT \rho(T)(1 - \beta e^{-T^*/T}). \tag{13}$$

Inserting the asymptotic behavior identified in Eq. (12) into Eq. (13) and assuming that the prefactor β does not depend on the characteristic failure time T yields the expected failure rate for large t,

$$\mathcal{F}(t) = 1 - 2\beta\sqrt{\alpha t} K_1(2\sqrt{\alpha t}), \tag{14}$$

where $K_1(x)$ is the first-order modified Bessel function of the second kind. The asymptotic expansion of this Bessel function for large arguments leads to

$$\mathcal{F}(t) \sim 1 - \beta\sqrt{\pi}(\alpha t)^{1/4} e^{-2(\alpha t)^{1/2}}, \tag{15}$$

which displays stretched exponential behavior. The analytic result of Eq. (14) provides a good fit to the time dependence of the failure probability, as shown in Fig. 23. For $\tau_D=0.18$, the fit with $\alpha=0.67\times10^{-5}$ and $\beta=0.91$ is good over essentially the whole time range. For $\tau_D=0.165$ (inset), the fit is not quite as good; the curvature of the analytic function seems to be smaller than that of the data. A fit to the $\tau_D=0.165$ data for $t\geq3400$ yields $\alpha=0.12\times10^{-3}$ and $\beta=0.64$. This fit (blue dotted line) underestimates the data for long times. A fit with $\alpha=0.10\times10^{-3}$ and $\beta=0.54$ (red dashed line) reduces the underestimation at long times, but is not as good for smaller times. In any case, the analytic results of Eqs. (14) and (15) indicate the existence of very long transients of prolonged activity. This activity will always eventually fail; $\mathcal{F}(t)\to1$ as $t\to\infty$.

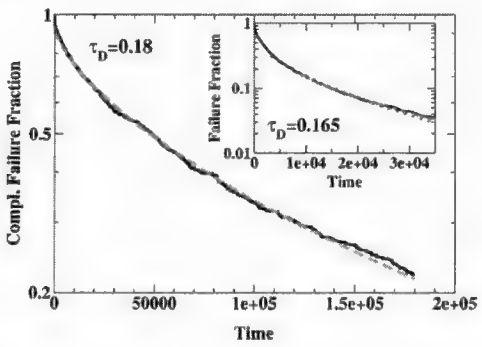

FIG. 23. (Color) Complementary failure fraction $1-\mathcal{F}$ as a function of time for $p=1$ and $\tau_D=0.18$ (inset for $\tau_D=0.165$). Dashed red and dotted blue lines are fits to Eq. (14). Note the logarithmic scale on the vertical axis.

V. CONCLUSION

In this paper we have used a minimal model to study the influence of network topology on the dynamics of coupled excitable elements. The network consists of a ring of locally connected elements; the connectivity is enhanced through random shortcuts that connect arbitrarily distant elements. Since our goal is to capture gross features of cortical connectivity, we assume that these shortcuts provide only unidirectional connections. This is in contrast to the bidirectional shortcuts that are appropriate for modeling epidemic propagation[19] or regular diffusive processes.

The dynamics of the system exhibit three distinct regimes, depending on the density of shortcuts and the speed of the waves that propagate through the network. For low but nonzero density of shortcuts, activity persists for essentially all network configurations when triggered by the initial excitation of a single neuron. This activity is predominantly periodic, and the mean firing rate of these persistent states shows only little dependence on the wave speed or the density of shortcuts once $p\geq0.05$. This firing rate is quite close to the maximal firing rate allowed by the recovery period of the neurons.

The recovery period is not to be confused with an absolute, intrinsic refractory period; rather, it is the time from one spike until the membrane potential has recovered to a value such that a single synaptic input of specified strength will suffice to trigger a new spike. This recovery period can be much longer than the absolute refractory period; this phenomenon has been observed in neurons that exhibit a slow after-hyperpolarization that underlies the slow oscillations (<1 Hz) observed *in vivo* in cat[31] and in cortical slices of ferret.[32] There, the recovery period induced by after-hyperpolarization can last as long as a few seconds. Of relevance to the dynamics investigated here is the time associated with the propagation speed of such slow oscillations over the whole network; specifically, the dependence of this propagation time on network connectivity. This dependence has been studied in cortical models[33,34] that do not incorporate shortcuts but use other mechanisms to control and vary network connectivity: either a variable spatial width of the Gaussian distribution that controls the probability that two neurons are connected,[34] or a trimodal probability distribution that captures a type of patchy connectivity in the cortex.[33] As expected, the speed of activity propagation was found to increase with increased connectivity; it was conjectured that this connectivity dependence underlies the large difference in propagation times observed for slow waves in olfactory cortex and neocortex.[34]

For low density of shortcuts, the persistent activity triggered by a localized excitation is periodic. However, the overall dynamical behavior of the system can be quite com-

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

plex due to the coexistence of a large number of stable solutions for a given network configuration. Whether the number of attractors grows as fast with system size as it does for globally coupled oscillator networks is not known at this point. In the case of global coupling, the permutation symmetry leads to a factorial growth of the number of attractors and to attractor crowding;[27] however, this permutation symmetry is broken in the small-world topology. So far as we have been able to investigate the dynamics of increasingly large networks, we have not been able to reach saturation of the number of attractors evoked through different initial conditions. The mechanism underlying this large number of attractors is not yet apparent; however, it is clear that noise will induce switching between these different attractors.[26]

As the density of shortcuts is increased, the number of network configurations that can sustain persistent activity decreases, until persistent activity is essentially no longer possible. For fast waves, the transients after a localized excitation consist of a single population burst followed by activity extinction. For slow waves, the crossover to complete failure occurs at larger shortcut densities than for fast waves. In addition, slow waves can exhibit exceedingly long transients that comprise thousands of population bursts. This activity eventually fails; the failure times for different network configurations at a fixed density of shortcuts follow an exponential distribution, which leads to a stretched exponential distribution for the expected probability of failure as a function of time. The mechanism responsible for these long transients is crucially dependent on the existence of pathways that bridge the quiescent periods between population bursts. These pathways are supported by the topology of the small-world network; this mechanism is thus expected to differ from the one relevant for dilute random networks of pulse-coupled oscillators.[35]

Another quantity of interest is the fraction of network configurations that has already failed by a specified time T^*. This quantity exhibits an overall decreasing trend with decreasing wave speed, or increasing recovery time τ_D. In addition to this decreasing trend, this quantity exhibits an intricate fine structure that includes sharp, resonance-like increases of the failure fraction with decreasing wave speeds. A naive argument supports the decreasing trend, but not the sharp resonances; a decrease in wave speed is likely to allow for additional shorter loops to contribute to the activity, and thus enhance the chances for persistence. However, the existence of sharp increases in failure probability indicates that the activation of one such loop can block the propagation of activity through previously active loops, and thus induce failure. While it is clear that such a switching between loops can occur with increasing τ_D, the mechanism that underlies these increased failures is not yet fully understood.

The crossover to failure, which can be understood in detail based on an analytic mean-field result for the effective size of these idealized small-world networks, provides the basis for understanding a number of recent studies of related but more complex neural network models.[4,22,23]

In Ref. 22, the connection between network connectivity and epilepsy in the hippocampus was investigated by considering small-world networks of three different types of neurons: noisy and leaky integrate-and-fire neurons, stochastic Hodgkin-Huxley cells, and Poisson spike-train cells. As parameters such as synaptic strength, number of synapses per neuron, and proportion of local versus shortcut connections, were varied, the network displayed dynamical behaviors described as "normal," "seizing," and "bursting." For low shortcut density, noise-driven activity was found to remain at a low level; this pattern of activity was associated with normal behavior. With increasing shortcut density, the level of activity triggered by a noise-driven event strongly increased, due to the recruitment of a vastly larger number of neurons. This pattern of activity was likened to seizing behavior. According to our analysis, this regime corresponds to patterns of connectivity that support persistent activity. Yet further increases in the shortcut density were found to induce bursting dynamics, characterized by irregular bursts involving a large fraction of all neurons, separated by quiescent periods. This pattern of activity corresponds in our analysis to network configurations associated with failure, for which each noise-triggered event leads to a population burst that brings essentially all neurons into their recovery period.

In Ref. 23, the role of network connectivity in sustaining population bursting activity was investigated by considering a small-world network of Morris-Lecar neurons. Activity was initiated through a localized set of pacemaker neurons. It was found that in the presence of shortcuts, network activity builds up over several driven cycles into bursts that involve a large fraction of neurons firing within a small time window. Both the time needed to build up such bursts and the time between them were found to decrease with increasing shortcut density. The appearance of repeated bursts is related to the failing configurations discussed here. As expected from our analysis, bursting behavior was supplanted by persistent activity when the wave speed was reduced (cf. increasing τ_D in Fig. 12). The slow buildup towards bursting activity observed in Ref. 23 appears to be specific to the Morris-Lecar neurons used in their model.

In Ref. 4, a 1:20 functional model of the rat dentate gyrus was constructed to investigate the functional consequences of two types of changes in network architecture that take place concurrently during epileptogenesis: loss of hillar mossy cells and sprouting of granular cell axons. The removal of mossy cells resulted in a massive reduction in the total number of connections, but this effect was to some extent compensated by an increase in local connectivity due to spatially restricted sprouting of granular cell axons. Surprisingly, as long as mossy cells were not almost fully eliminated, they continued to provide intermediate and long-range connections onto granule cells, thus preserving the short network diameter characteristic of small-world connectivity. The hyperexcitability that arises from these structural changes corresponds in our analysis to the sustained propagation of fast waves in the presence of a small density of shortcuts.

ACKNOWLEDGMENTS

We gratefully acknowledge the NSF support through Grant Nos. DMS-0309657 and DMS-322807, as well as the IGERT program "Dynamics of Complex Systems in Science

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

and Engineering" (DGE-9987577) (H.R., A.R., S.M.); EU support under Grant No. MRTN-CT-2004-005728 (S.M.) is also gratefully acknowledged. A.R. acknowledges funding by Marie Curie grant 022242.

[1]R. Albert and A. L. Barabasi, Rev. Mod. Phys. **74**, 47 (2002).
[2]M. E. Newman, SIAM Rev. **45**, 167 (2003).
[3]O. Shefi, I. Golding, R. Segev, E. Ben-Jacob, and A. Ayali, Phys. Rev. E **66**, 021905 (2002).
[4]J. Dyhrfjeld-Johnsen, V. Santhakumar, R. J. Morgan, R. Huerta, L. Tsimring, and I. Soltesz, J. Neurophysiol. **97**, 1566 (2007).
[5]M. Barahona and L. M. Pecora, Phys. Rev. Lett. **89**, 054101 (2002).
[6]H. Hong, M. Y. Choi, and B. J. Kim, Phys. Rev. E **65**, 026139 (2002).
[7]T. Nishikawa, A. E. Motter, Y.-C. Lai, and F. C. Hoppensteadt, Phys. Rev. Lett. **91**, 014101 (2003).
[8]P. C. Bressloff, J. Math. Biol. **40**, 169 (2000).
[9]P. C. Bressloff, Physica D **155**, 83 (2001).
[10]D. Golomb and G. B. Ermentrout, Phys. Rev. Lett. **86**, 4179 (2001).
[11]R. D. Traub, D. Schmitz, J. G. Jefferys, and A. Draguhn, Neuroscience **92**, 407 (1999).
[12]T. J. Lewis and J. Rinzel, Network Comput. Neural Syst. **11**, 299 (2000).
[13]T. J. Lewis and J. Rinzel, Neurocomputing **38**, 763 (2001).
[14]M. Perc, New J. Phys. **7**, 252 (2005).
[15]L. F. Lago-Fernández, R. Huerta, F. Corbacho, and J. A. Sigüenza, Phys. Rev. Lett. **84**, 2758 (2000).
[16]L. F. Lago-Fernández, F. J. Corbacho, and R. Huerta, Neural Networks **14**, 687 (2001).
[17]D. R. Paula, A. D. Araujo, J. S. Andrade, H. J. Herrmann, and J. A. Gallas, Phys. Rev. E **74**, 017102 (2006).
[18]A. R. Carvunis, M. Latapy, A. Lesne, C. Magnien, and L. Pezard, Physica A **367**, 595 (2006).
[19]M. Kuperman and G. Abramson, Phys. Rev. Lett. **86**, 2909 (2001).
[20]A. Roxin, H. Riecke, and S. A. Solla, Phys. Rev. Lett. **92**, 198101 (2004).
[21]D. J. Watts and S. H. Strogatz, Nature (London) **393**, 440 (1998).
[22]T. I. Netoff, R. Clewley, S. Arno, T. Keck, and J. A. White, J. Neurosci. **24**, 8075 (2004).
[23]J. Shao, T. Tsao, and R. Butera, Neural Comput. **18**, 2029 (2006).
[24]M. Tinsley, J. X. Cui, F. V. Chirila, A. Taylor, S. Zhong, and K. Showalter, Phys. Rev. Lett. **95**, 038306 (2005).
[25]A. J. Steele, M. Tinsley, and K. Showalter, Chaos **16**, 015110 (2006).
[26]A. Roxin, Ph.D. thesis, Northwestern University (2003).
[27]K. Wiesenfeld and P. Hadley, Phys. Rev. Lett. **62**, 1335 (1989).
[28]M. E. Newman, C. Moore, and D. J. Watts, Phys. Rev. Lett. **84**, 3201 (2000).
[29]M. E. Newman and D. J. Watts, Phys. Rev. E **60**, 7332 (1999).
[30]T. Petermann and P. de Los Rios, Phys. Rev. E **73**, 026114 (2006).
[31]F. Amzica and M. Steriade, J. Neurophysiol. **73**, 20 (1995).
[32]M. V. Sanchez-Vives and D. A. McCormick, Nat. Neurosci. **3**, 1027 (2000).
[33]A. Compte, M. V. Sanchez-Vives, D. A. Mccormick, and X. J. Wang, J. Neurophysiol. **89**, 2707 (2003).
[34]M. Sanchez-Vives and A. Compte, Lect. Notes Comput. Sci. **3561**, 133 (2005).
[35]A. Zumdieck, M. Timme, T. Geisel, and F. Wolf, Phys. Rev. Lett. **93**, 244103 (2004).

Downloaded 05 Oct 2007 to 130.207.50.192. Redistribution subject to AIP license or copyright, see http://chaos.aip.org/chaos/copyright.jsp

Karel Svoboda – HHMI, Janelia Farm Research Campus
svobodak@janelia.hhmi.org
http://research.janelia.org/labs/display/SvobodaLab/Home

1. Synaptic transmission

- Basics – presynaptic / postsynaptic; Inhibitory / excitatory; ionotropic receptors
- Spatial aspects of synaptic transmission – synapse-specificity
- Noise– synaptic noise dominates other neuronal noise sources
- Linearity
- Speed
- Structure-function relationships

2. Synaptic plasticity

Useful references

- Hestrin et al J. Physiol 1990 – classic analysis of synaptic receptors in brain slices
- Tsodyks and Markram 1997; Abbott, L, et al 1996 – classic papers on the effects of short-term depression on signal propagation in neural networks
- Regehr and Zucker 2002 – excellent review on short-term plasticity
- Sabatini and Regehr 1999 – only recent review on timing in synaptic transmission
- Chklovskii, Mel, Svoboda 2004 – structural plasticity
- Dan and Poo 2006 – review of STDP
- Harvey and Svoboda 2007 – coupling between synapses
- And about 124,993 more ….

Synapses and circuits

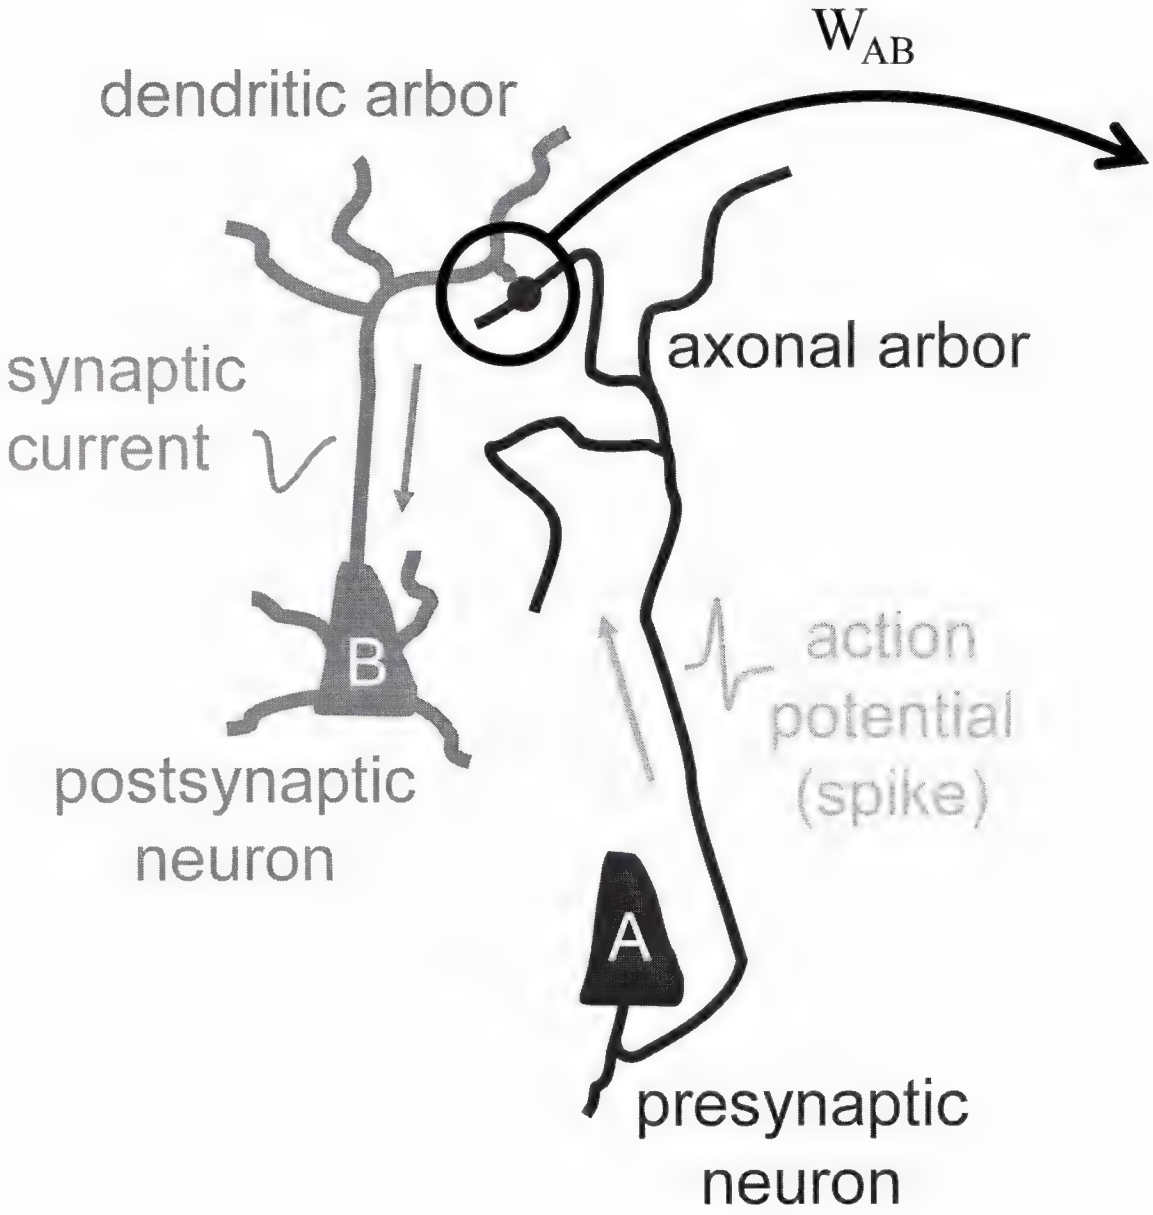

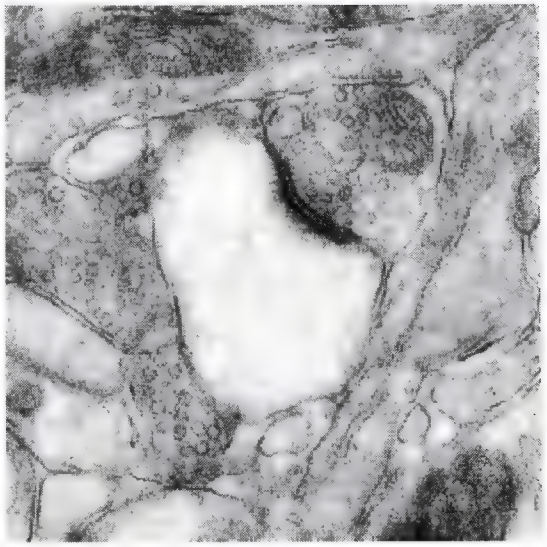

0.5 μm

Excitatory and inhibitory synapses

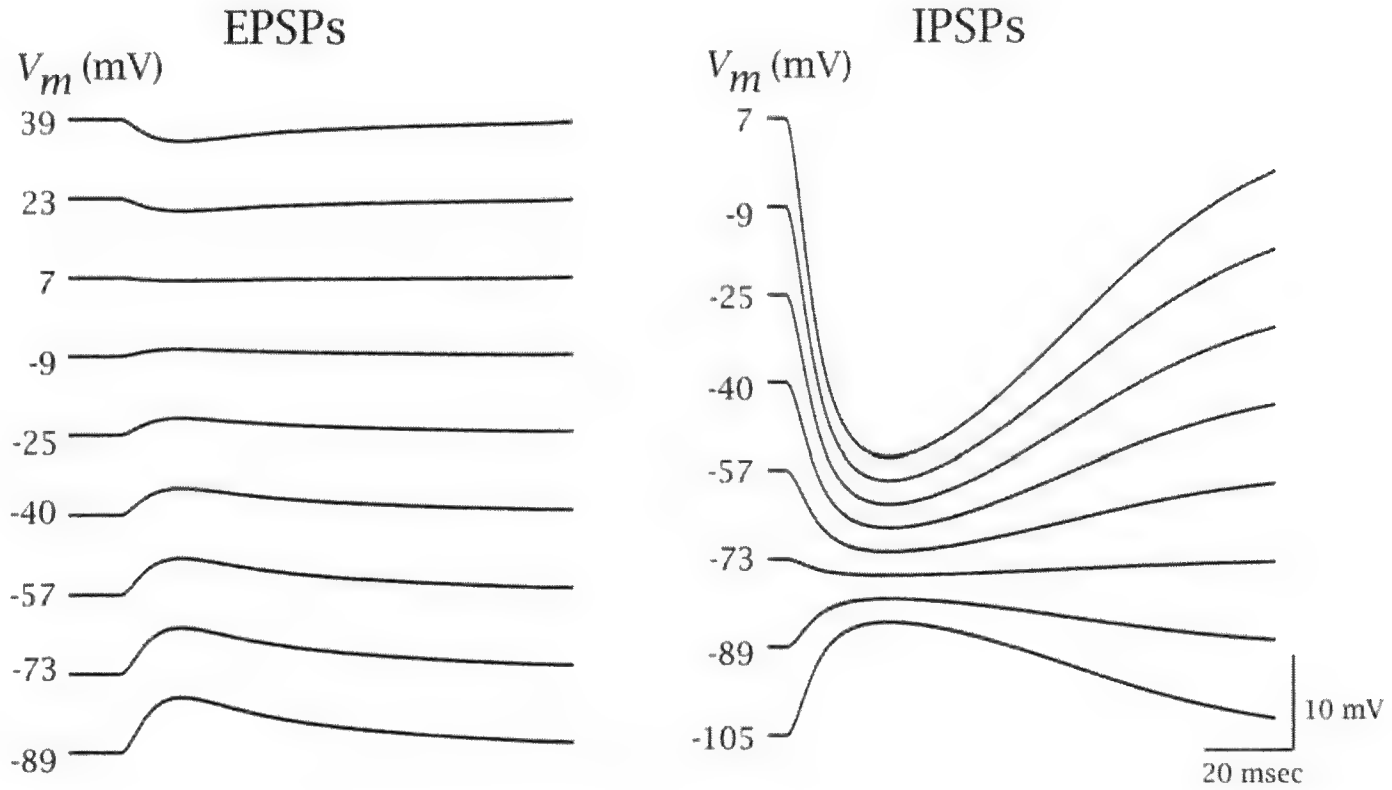

Figure 13.4 Examples of EPSPs and IPSPs at different membrane potentials in a passive neuron.

Measuring synaptic transmission

Hippocampal brain slice

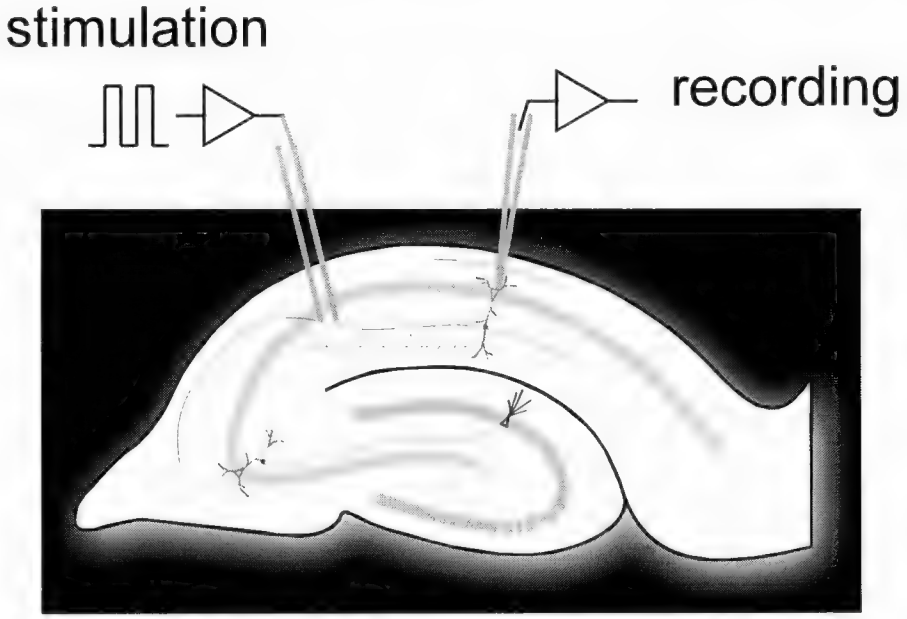

Excitatory postsynaptic current, EPS**C**

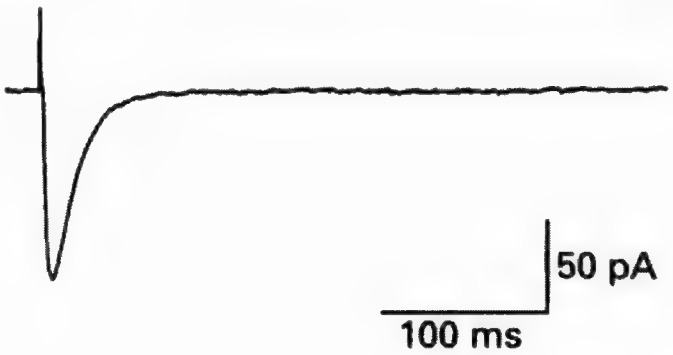

Reresents function of 10-20 synapses

Distributed throughout the dendrite

Hestrin et al 1990

Two major types of ionotropic glutamate receptors

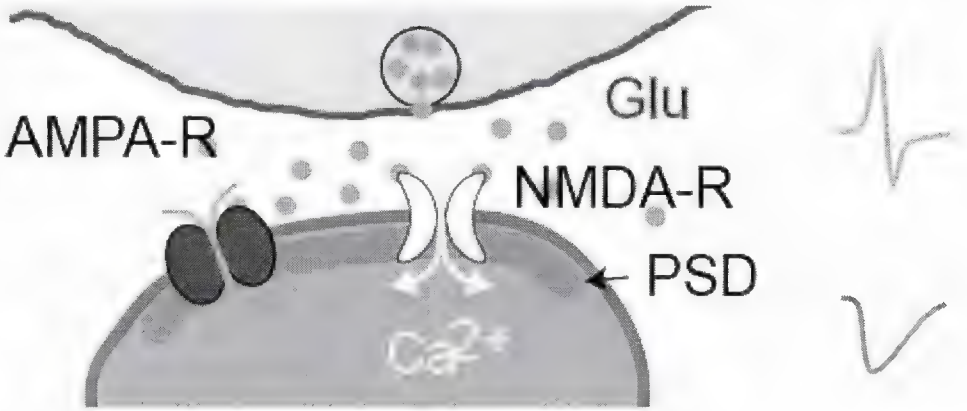

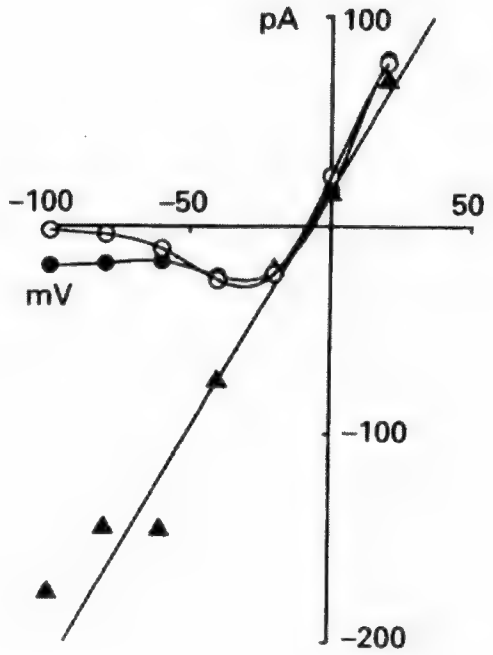

Block
NMDA-receptors

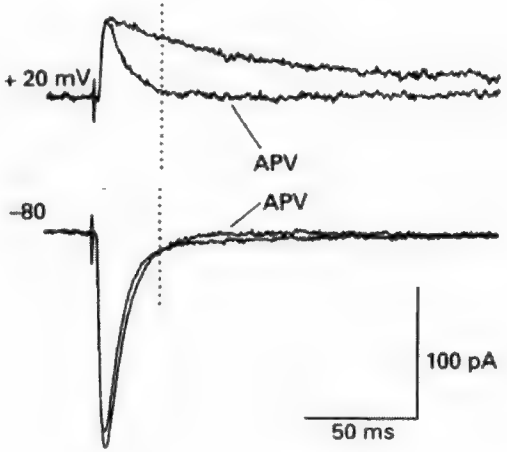

Block
AMPA-receptors

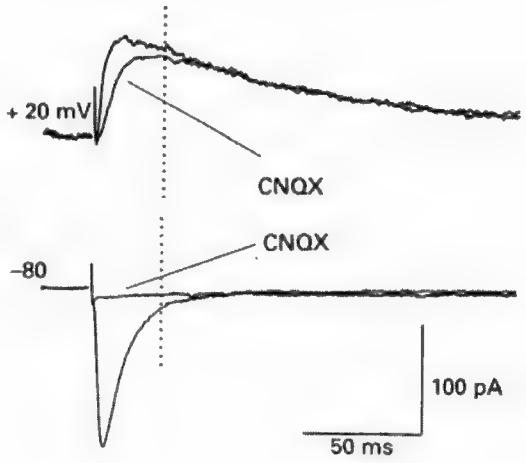

Hestrin et al 1990

Imaging to study individual synapses in their habitat

High-resolution, non-invasive – image individual synapses for months

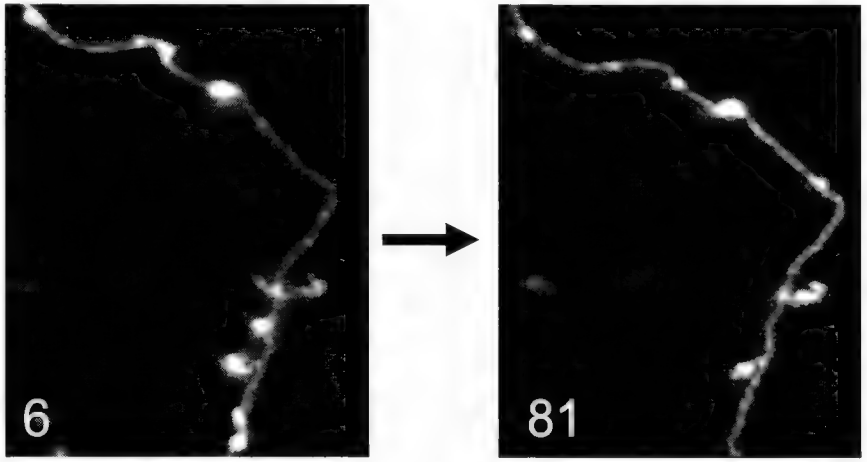

Versatile - probe aspects of function (e.g. Ca^{2+})

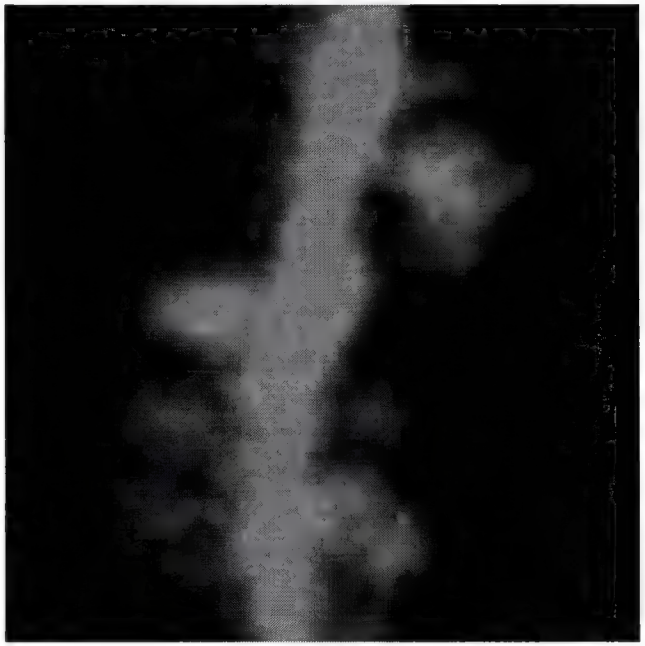

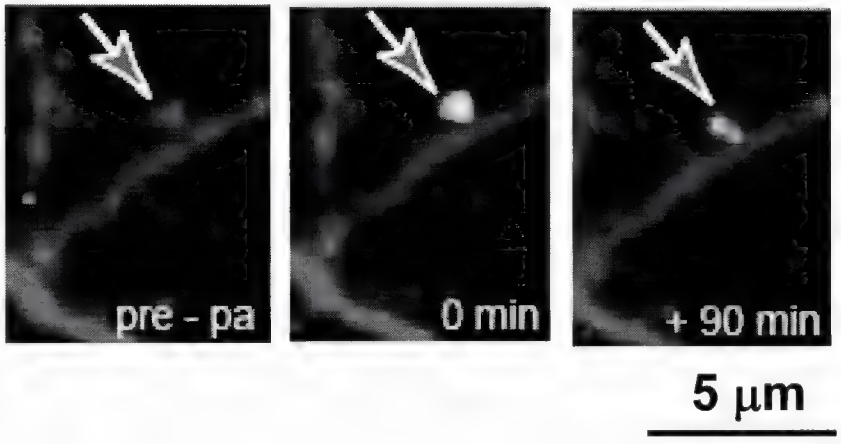

Specific - measure locations & trafficking of molecules

Imaging synaptic transmission and Ca^{2+} signaling

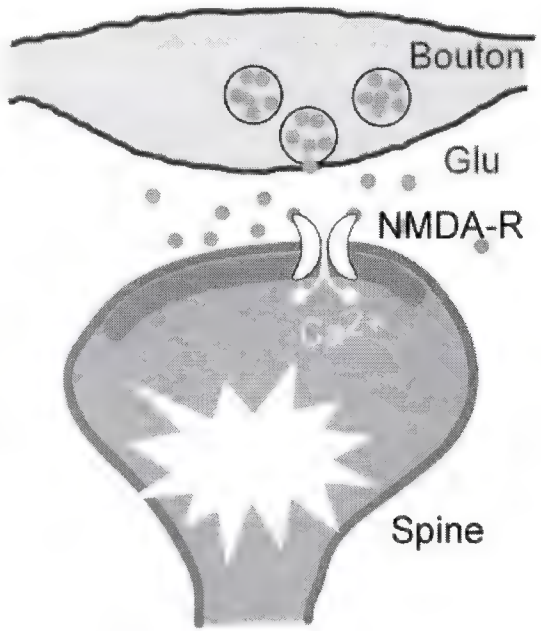

Imaging Ca^{2+} through NMDA –Rs:

1. Sensitive method to detect glutamate
2. Optical measurement of glutamate release at single synapses

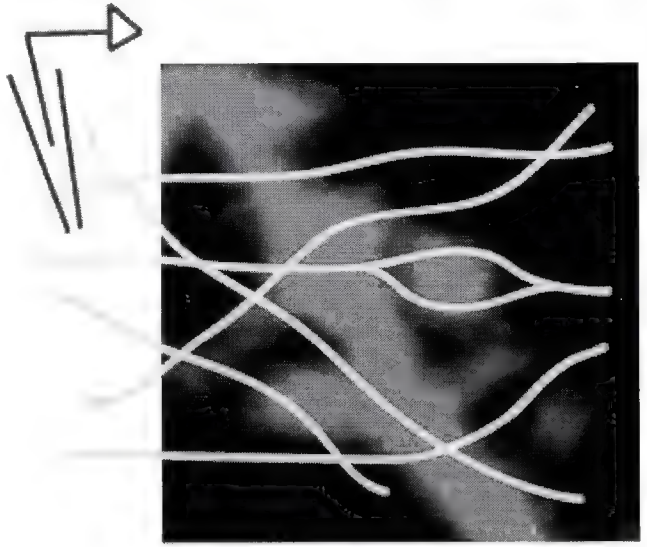

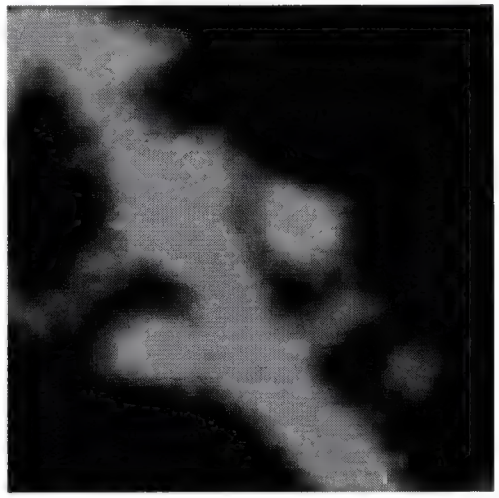

- **Synaptic transmission is point-to-point**
- **Few receptors activated -> noise**
- **Spine [Ca2+] signals are synapse-specific**

Yuste and Denk, 1995
Mainen et al 2000
Oertner et al 2002

NMDA-Rs are not saturated by release due to a single action potential

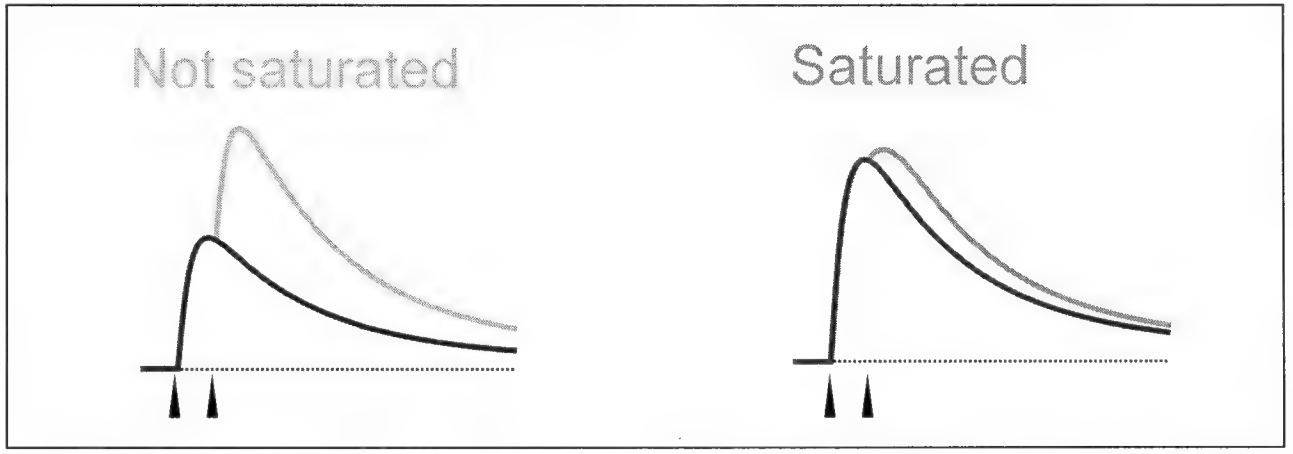

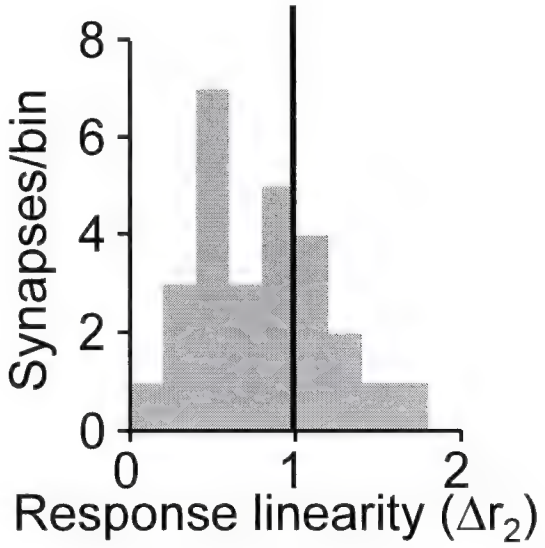

Average
$\Delta r_2 = 0.8 \pm 0.08 r_1$;

2. stim in a pair gives > 80% additional activation

Mainen, Malinow, KS. 1999. Nature

•NMDA-R activation reports the amount of glutamate released

Glutamate receptors are not saturated!

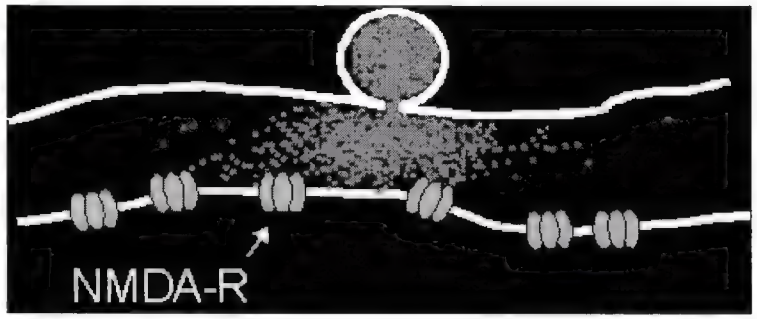

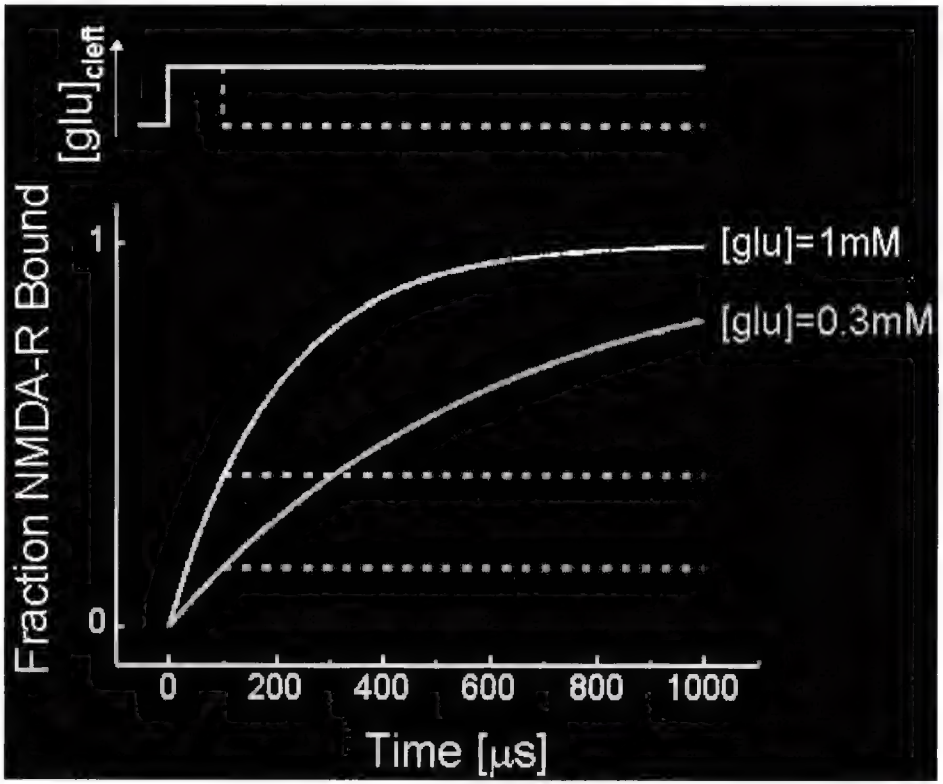

cleft volume ~$1x10^{-18}$ liter
# of glu/vesicle ~ 1,000
# of receptors ~ 10
peak [glu] ~ 1mM
Kd (NMDA-R) ~ 1 μM

k_{on} ~ $5x10^6$/Ms
k_{off} ~ 5/s

Conundrum solved: glutamate transients are very brief;
extremely rapid uptake of glutamate from the extracellular space

Synaptic transmission is extremely fast

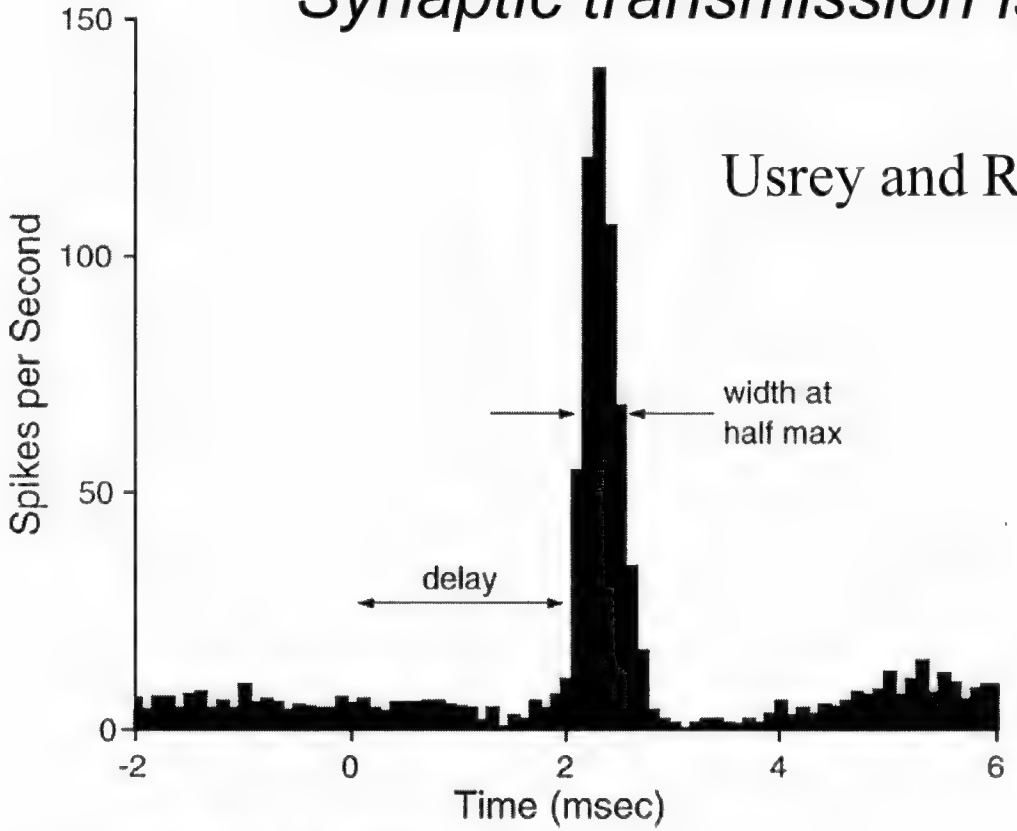

Usrey and Reid

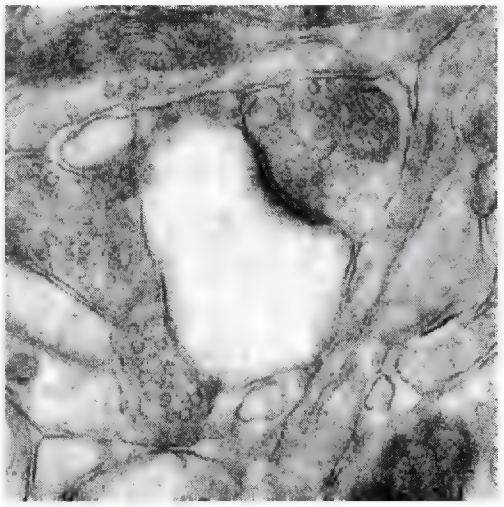

0.5 μm

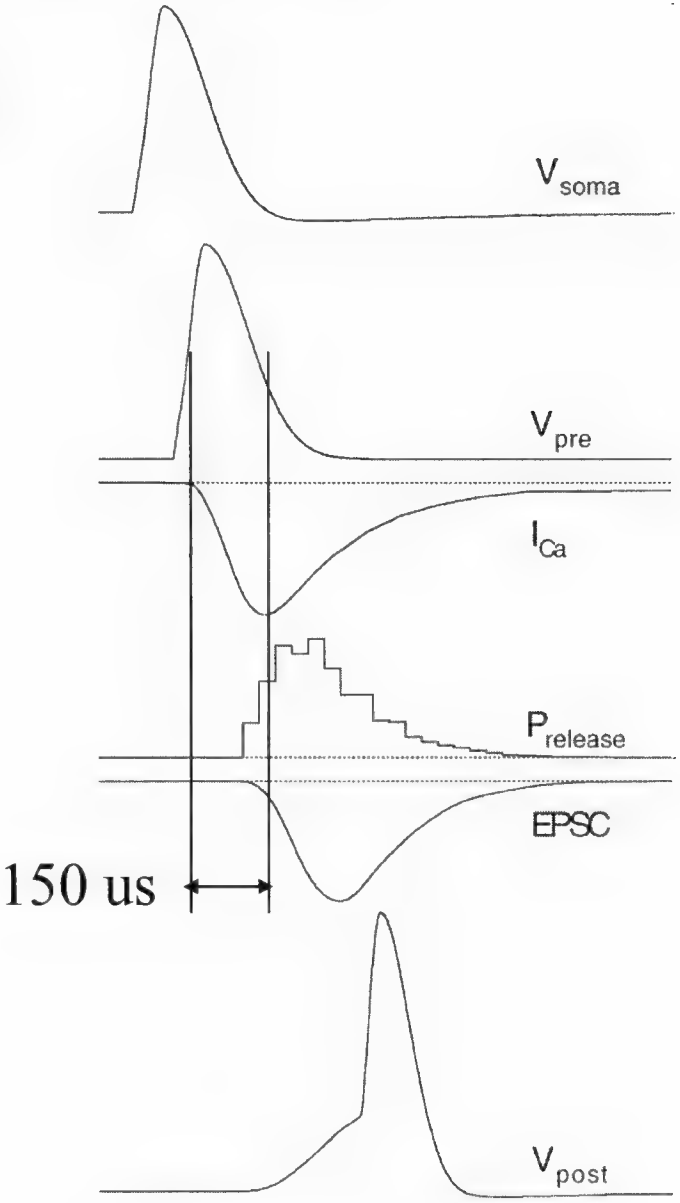

Sabatini and Regehr

1. Synaptic transmission

- Input specific; point-to-point
- Receptors are not saturated; postsynaptic response is linear
- NT-release is graded at single synapses …
- Single cortical synapses are noisy
- Extremely fast

Synaptic plasticity

Short-term – (milliseconds to seconds)
'Computation'

Long-term – (tens of minutes to years)
'Learning and memory'

Functional and structural
Induction protocols -- STDP
Structural
Synapse-specificity
Crosstalk between plasticity at nearby synapses and its mechanisms

Homeostatic plasticity

Short-term synaptic plasticity -facilitation

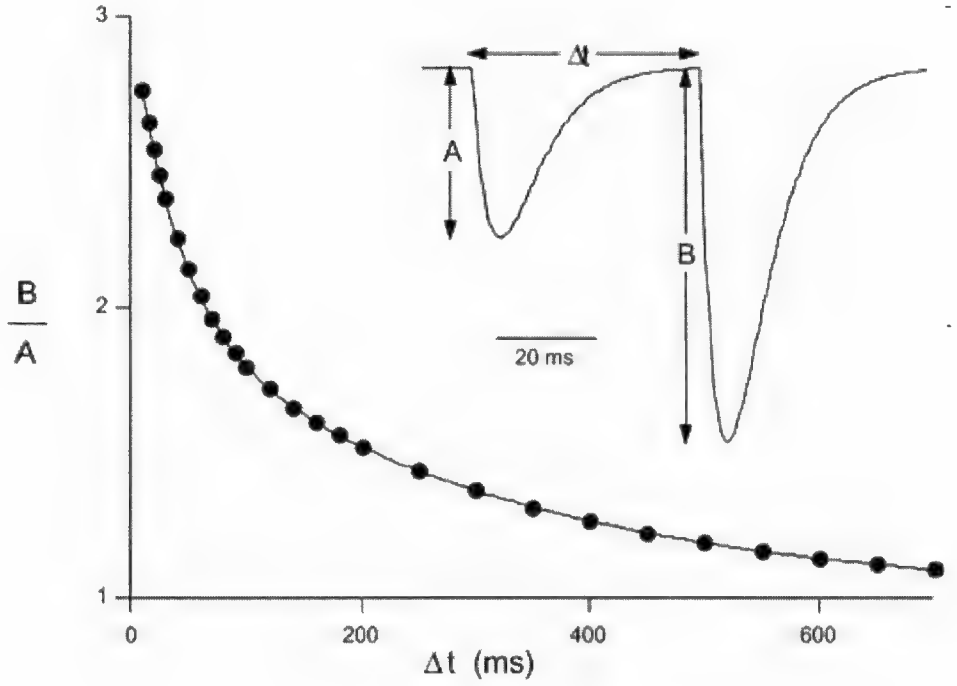

Zucker and Regher

Dodge-Rahamimoff Relation

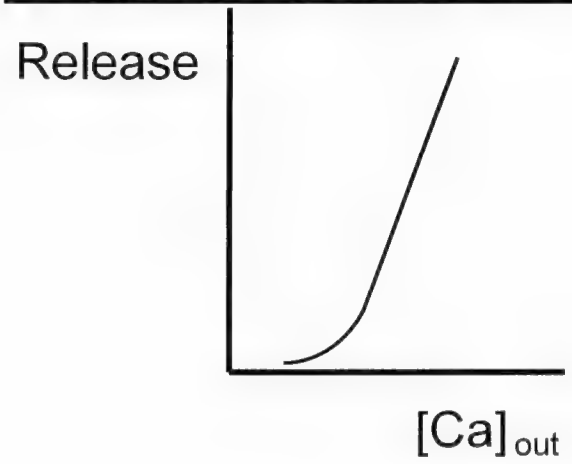

Cooperativity

$$R = [Ca_o]^n$$

- Strictly feedforward …
- Mechs. of short-term plasticity are mainly presynaptic
- Facilitation is in part due to residual calcium (NMJ; Miledi and Katz)

Many time-scales and mechanisms of facilitation and depression

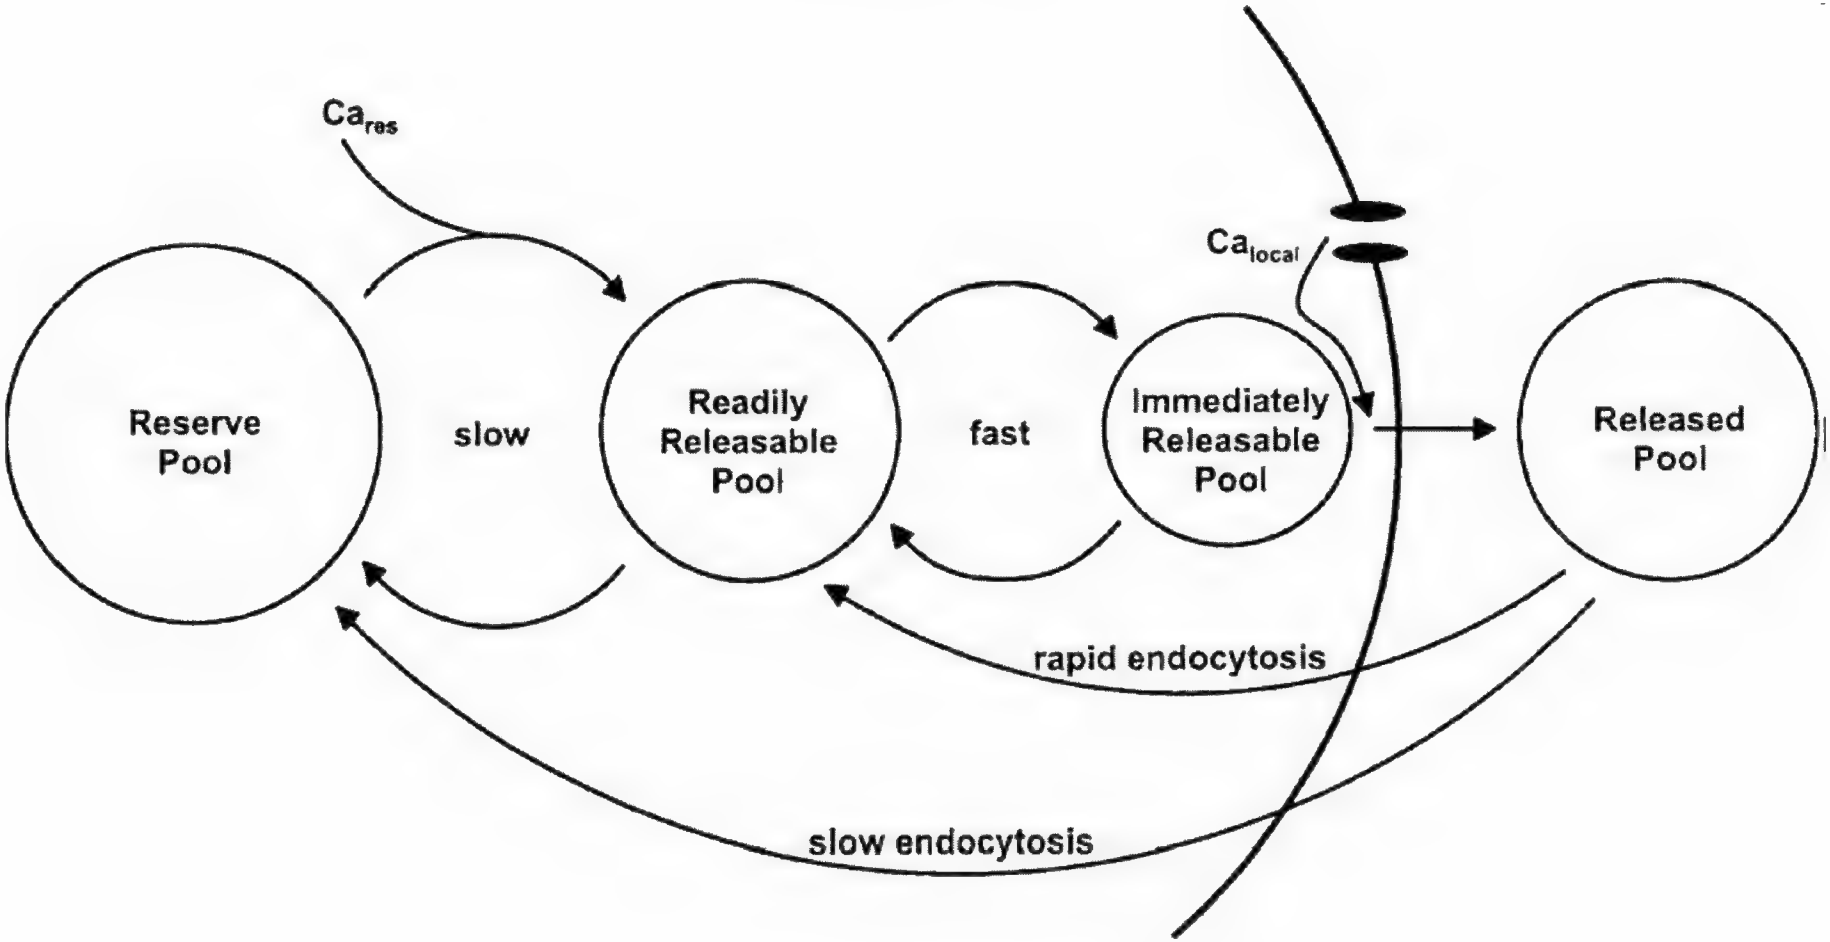

Depression at cortical synapses

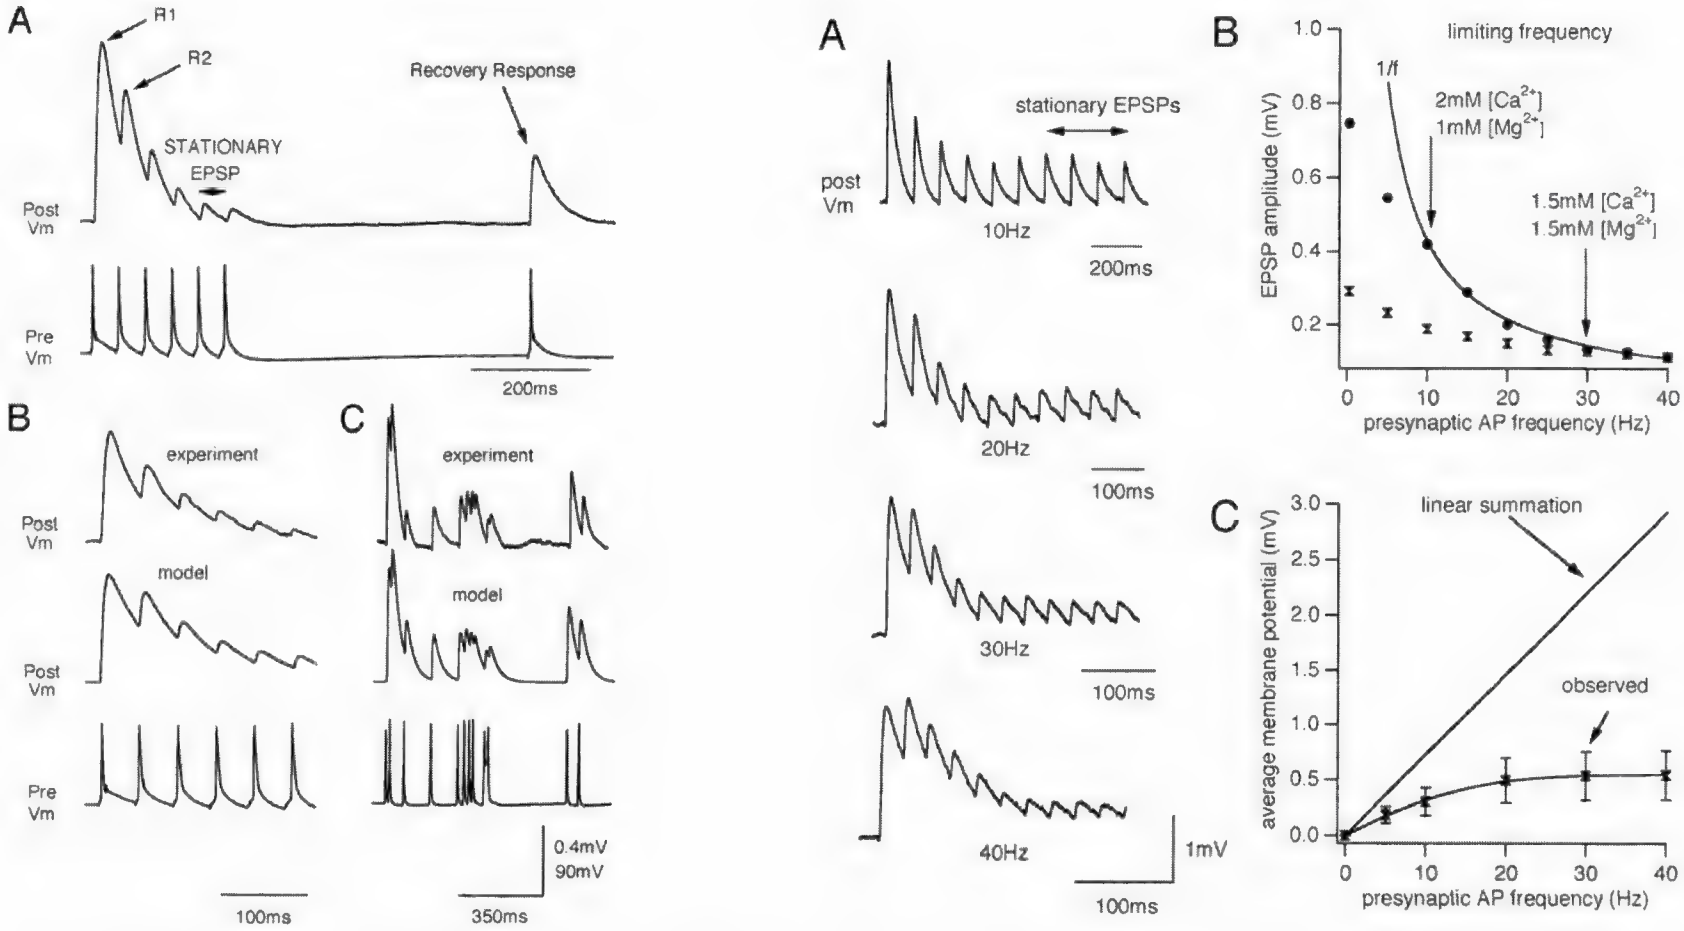

Tsodyks M, Markram H PNAS 1997;94:719-723; also Abbott et al 1996

Short-term synaptic depression – depletion of a pool of vesicles (docked vesicles) that are poised and ready for release

Depression at cortical synapses - differentiation

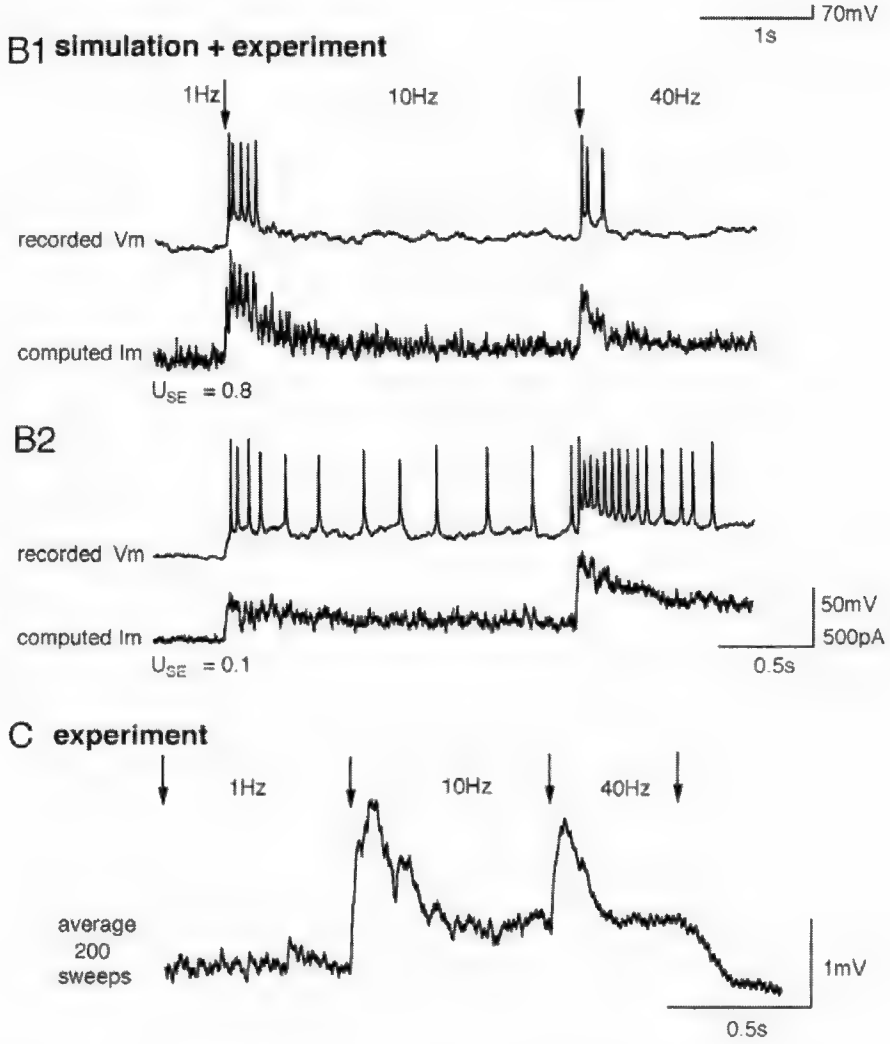

Tsodyks M, Markram H PNAS 1997;94:719-723

See also Abbott, L et al 1996

Mechanisms of paired-pulse facilitation (PPF)

syn. current measured in soma

number of docked vesicles, ***N***
vesicular rel. prob, ***p***

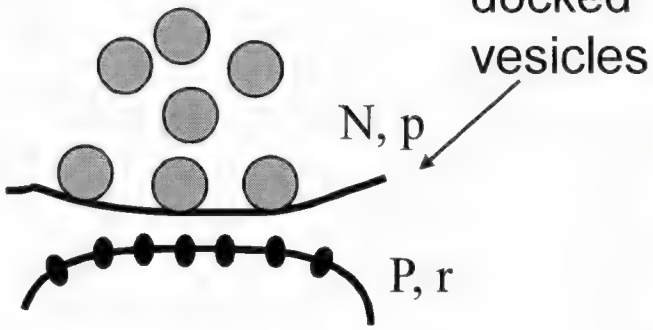

potency ***r***
syn. rel. prob. ***P***

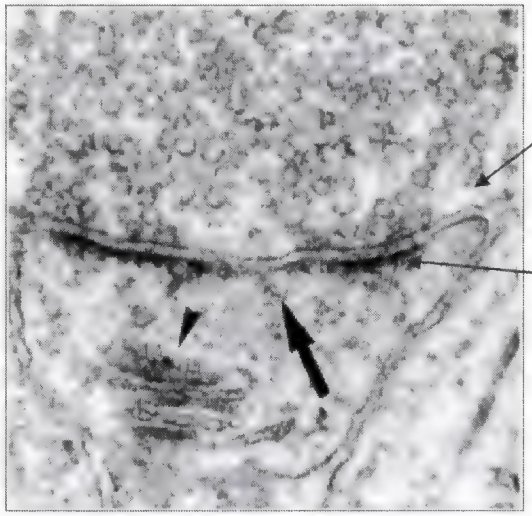

Schikorski & Stevens

Mechanisms of paired-pulse facilitation (PPF)

Independent release sites

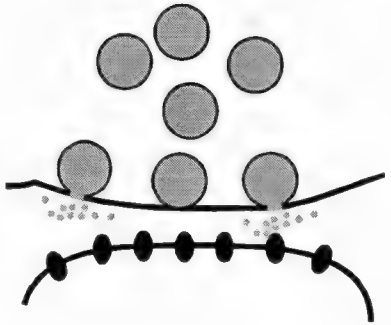

- *potency, **r**, variable*
- ***r** dependent on release prob.*
- *PPF due to mod. of **P** and **r***
- *'multivesicular' release*

$$\frac{r_2}{r_1} = \frac{P_1\left(1 - P_2\right)^{1/N_d}}{P_2\left(1 - P_1\right)^{1/N_d}}$$

Univesicular dogma

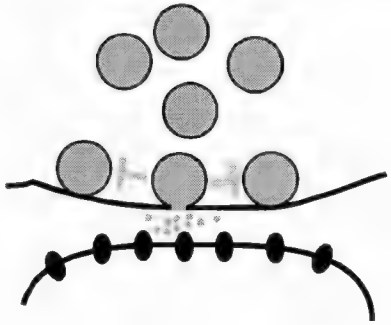

- *potency, **r**, constant*
- ***r** independent of release prob.*
- *PPF due to modulation of **P** alone*
- *'univesicular' release*

$$\frac{r_2}{r_1} = 1$$

- $[Ca^{2+}]$ imaging in spines allows us to measure both ***P*** and ***r*** at single synapses -- test models of plasticity

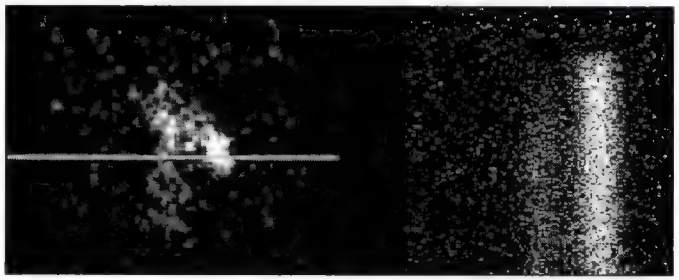

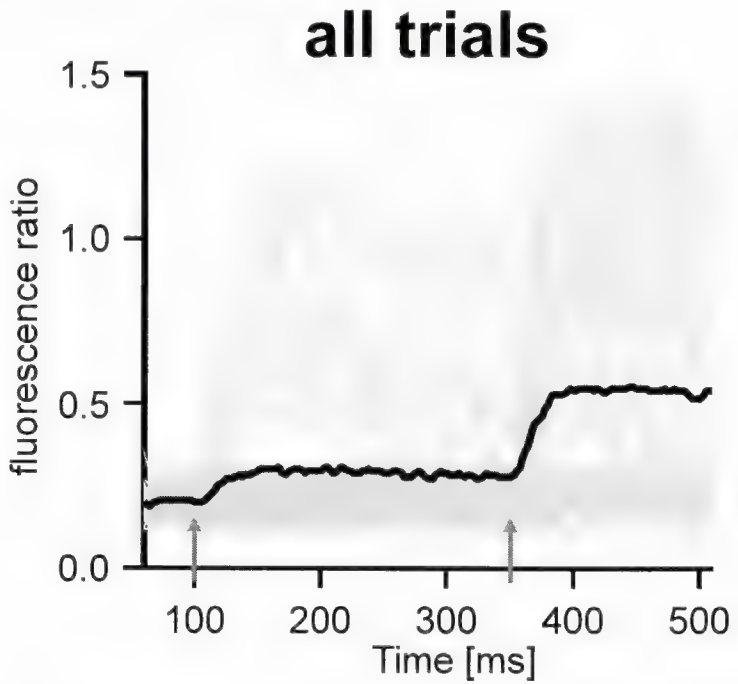

Sorting into outcomes

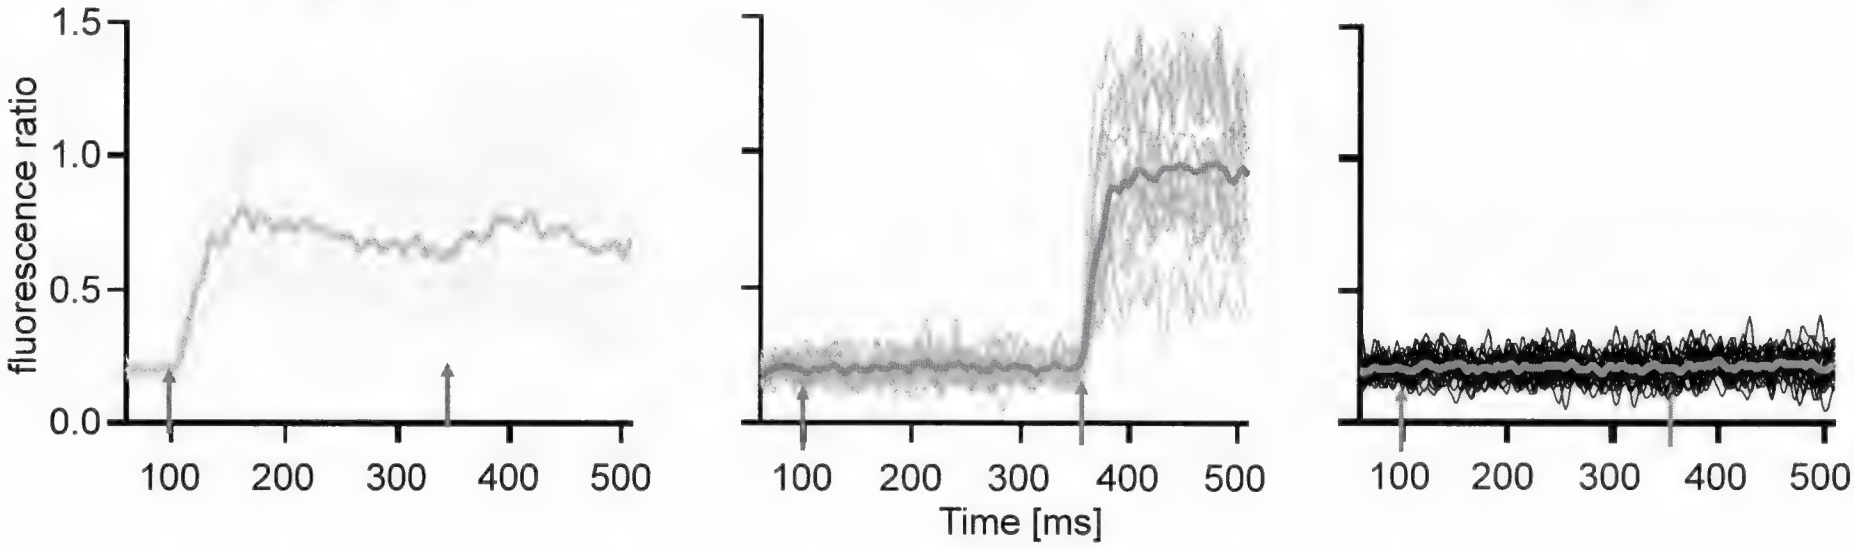

Paired pulse facilitation dissected:

*both **p** and **r** are modulated*

Change in ***P***

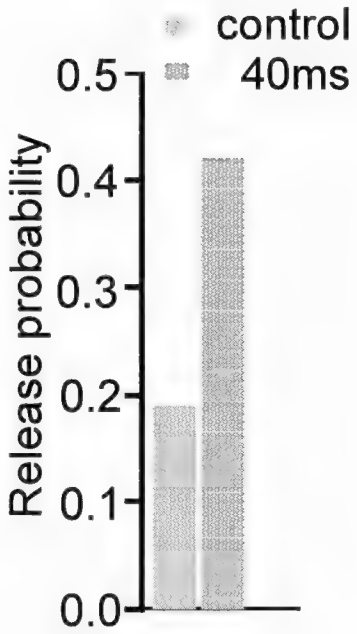

Change in potency, ***r***

Successes only,

given failure to 1st pulse

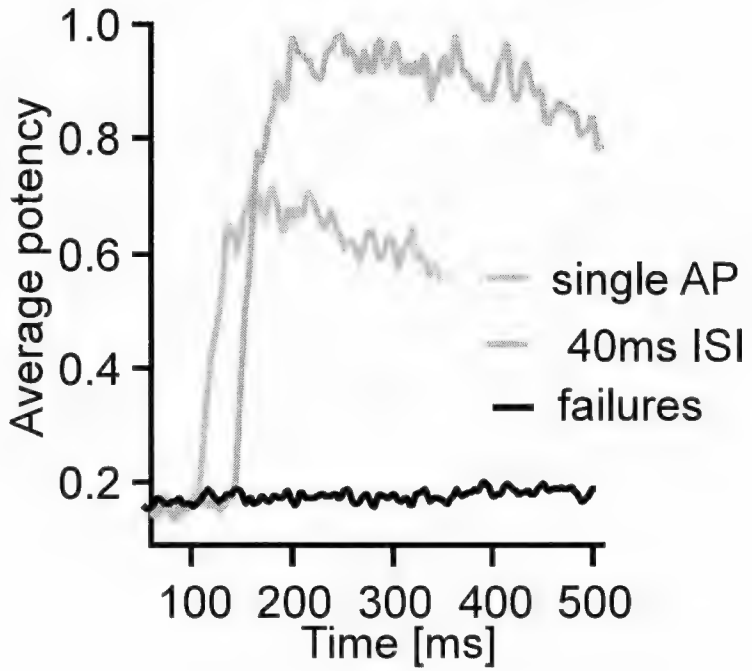

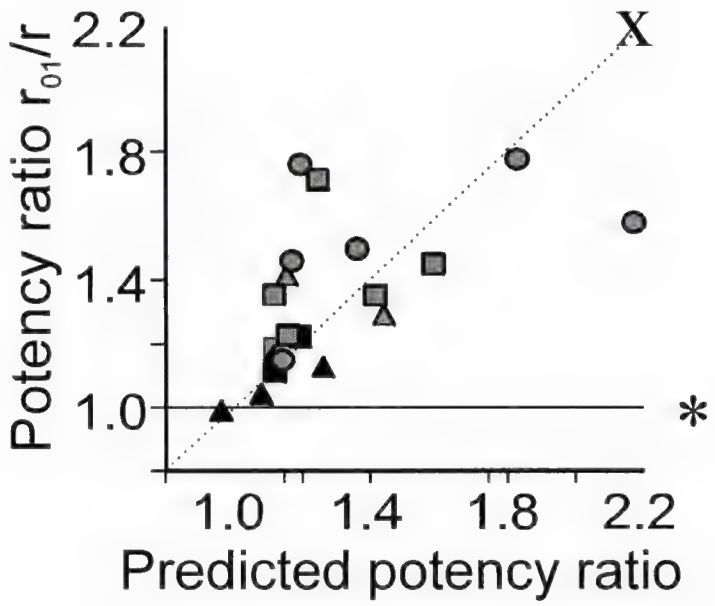

News

Nature **449**, 124-125 (13 September 2007) | doi:10.1038/449124b; Published online 12 September 2007

Long-held theory is in danger of losing its nerve

$$\frac{r_{01}}{r_1} = \frac{P_1(1-P_2)^{1/N_d}}{P_2(1-P_1)^{1/N_d}} \quad \text{X} \qquad N_d = 5$$

$$\frac{r_{01}}{r_1} = 1 \quad *$$

Oertner et al 2002; Christie and Jahr 2006

Imaging structural plasticity in vivo

Two-photon laser scanning microscope

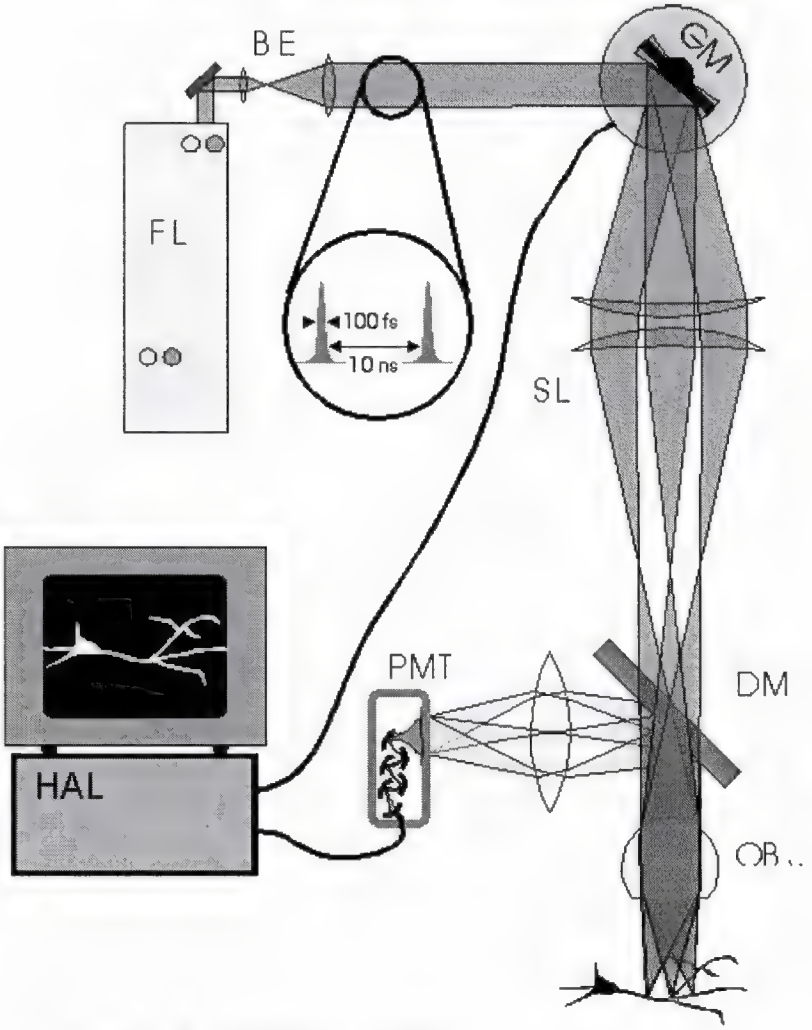

Denk & Webb, 1990

Reviewed in Svoboda & Yasuda, 2006

Intact mouse

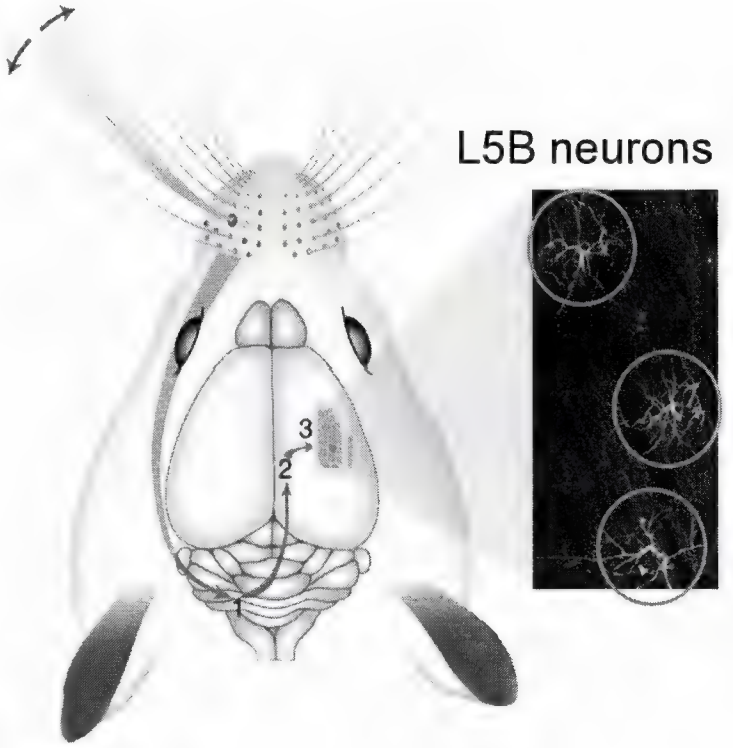

New dendritic spines after novel experience

Long-term in vivo imaging

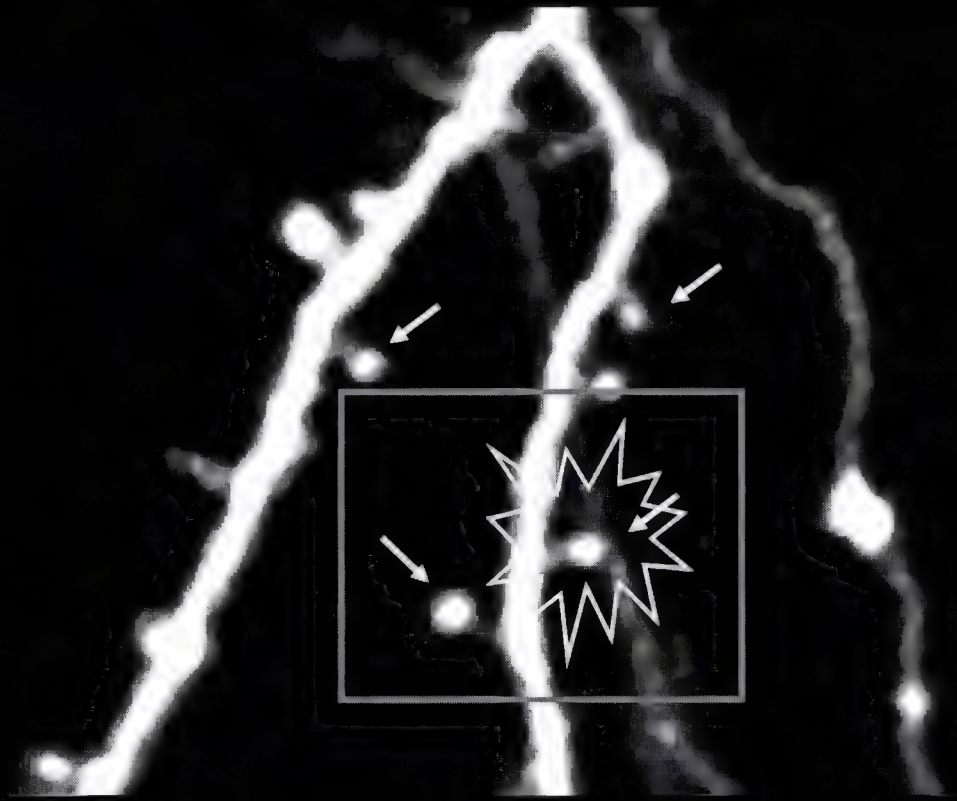

5 μm

Experimental time-course

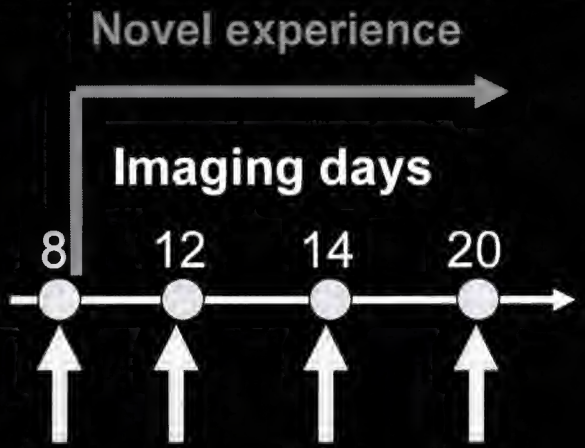

- A subpopulation of spines appear and disappear over days
- New spines are selectively stabilized after novel experience

New spines have synapses

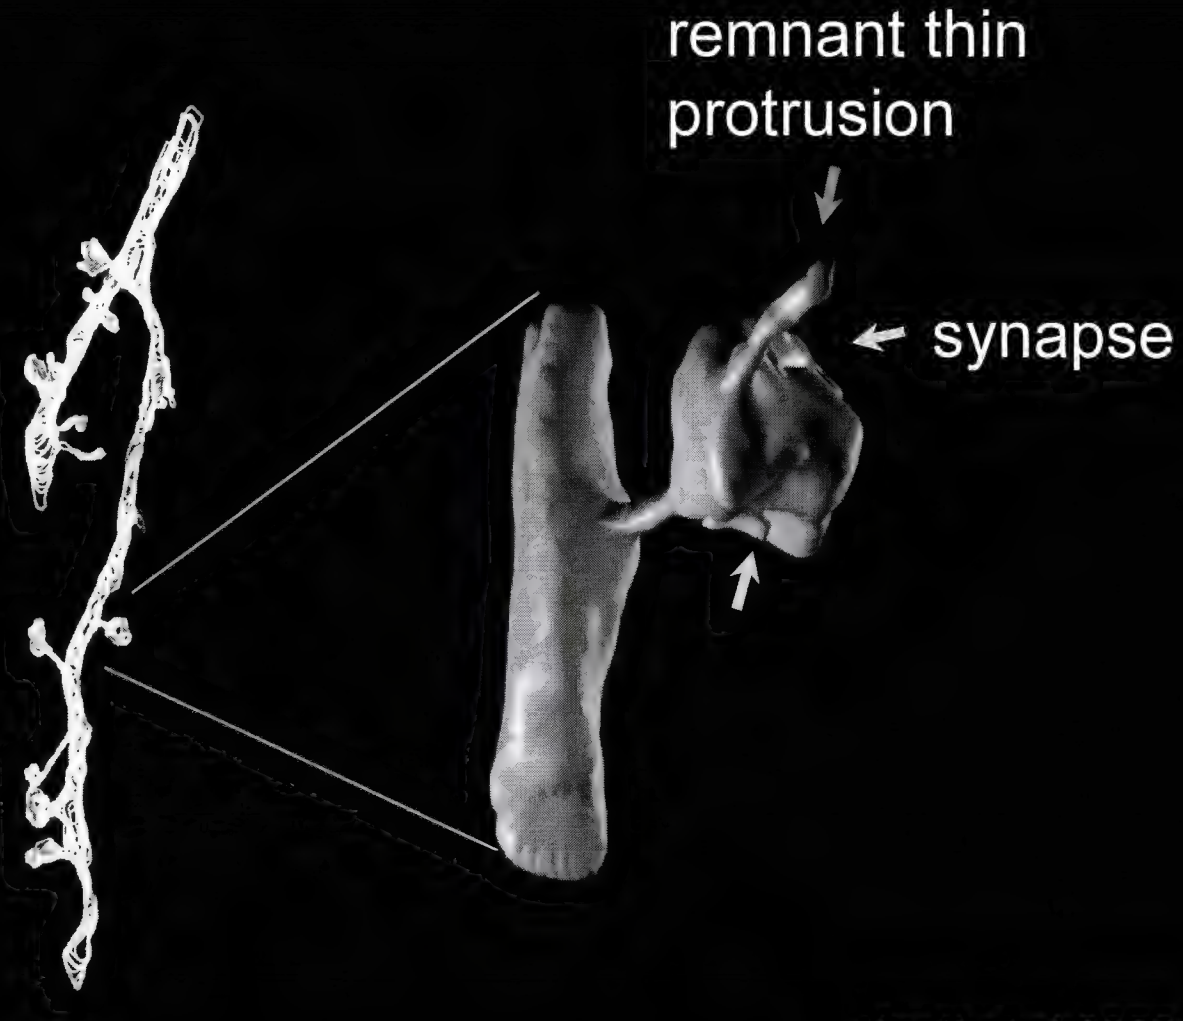

Knott et al 2006
Holtmaat et al 2006

'Sampling' and experience-dependent 'selection'

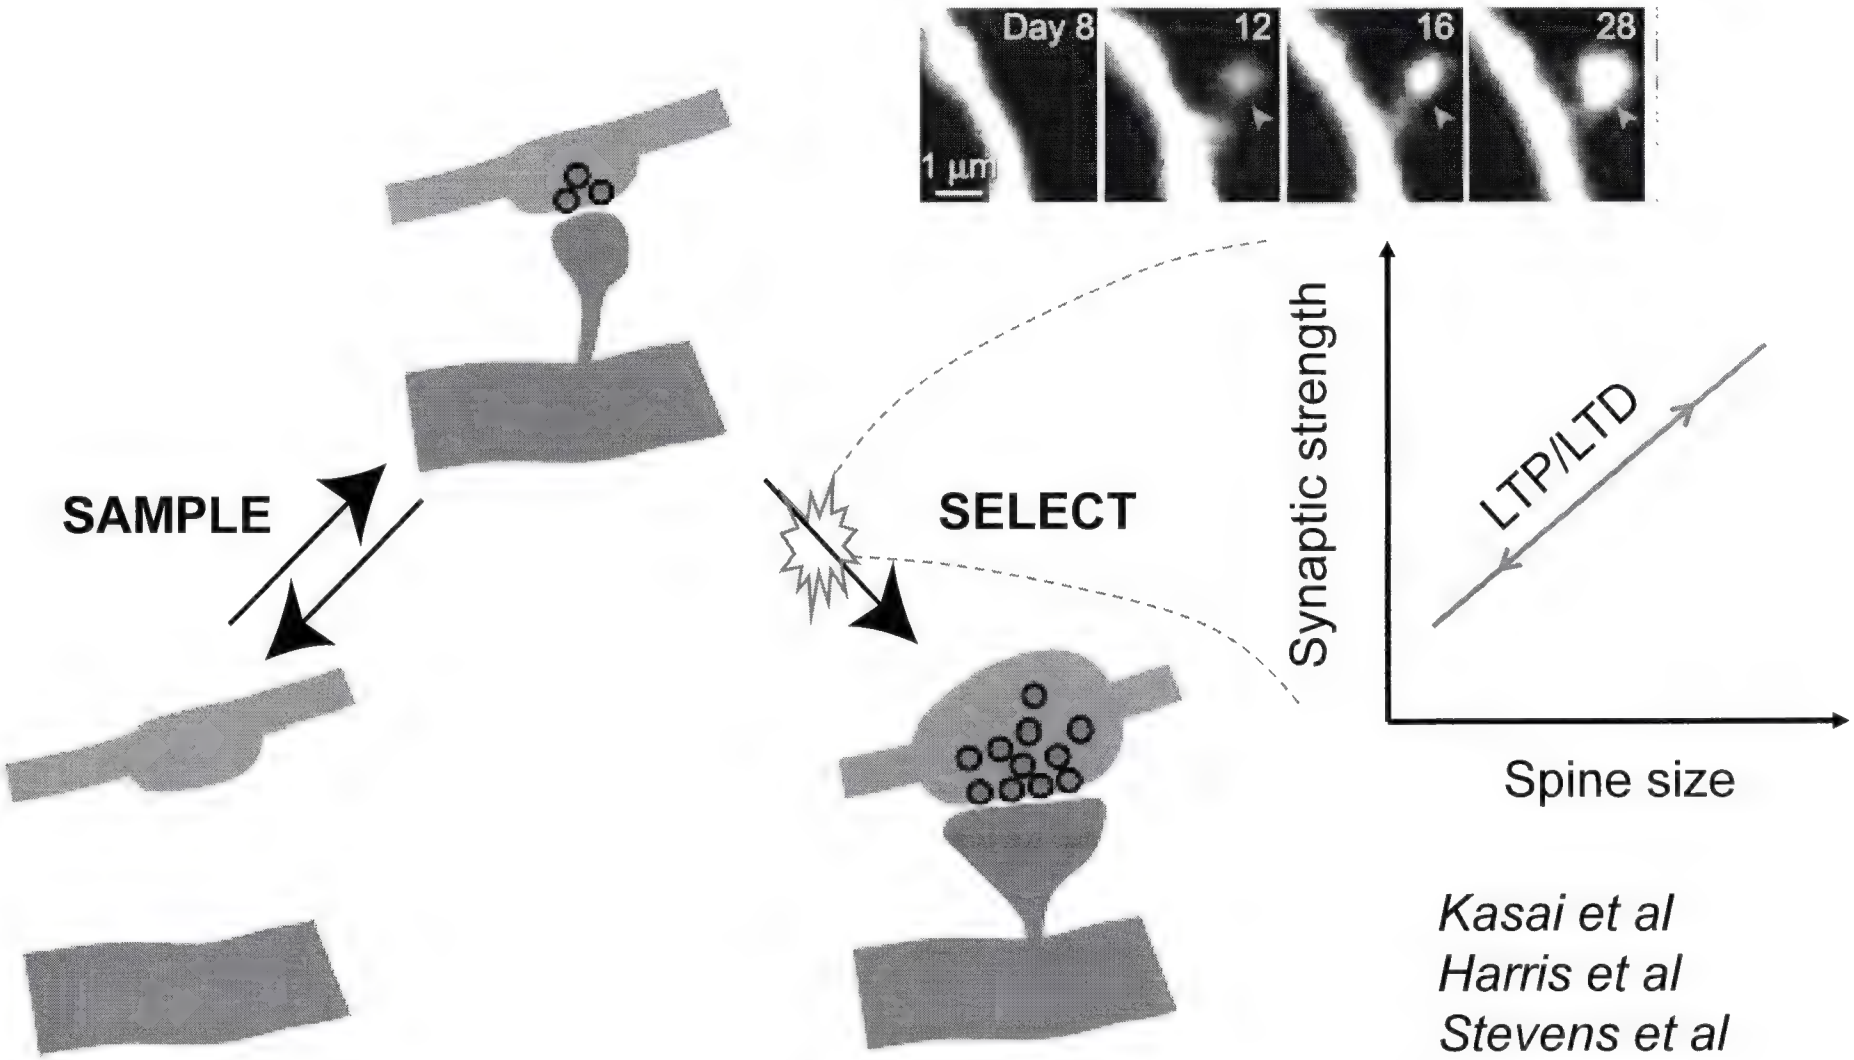

Kasai et al
Harris et al
Stevens et al

Trachtenberg et al 2002; Holtmaat et al 2005; Knott et al 2006; Holtmaat et al 2006.
Wilbrecht et al (unpublished); Kuhlman et al (unpublished)

Why have structural plasticity ?

- Storage capacity

$\sim -Np\,[\,(1-f)\log_2(1-f) + f\log_2 f\,]$ (bits)

- Np – number of potential synapses
- f – filling fraction

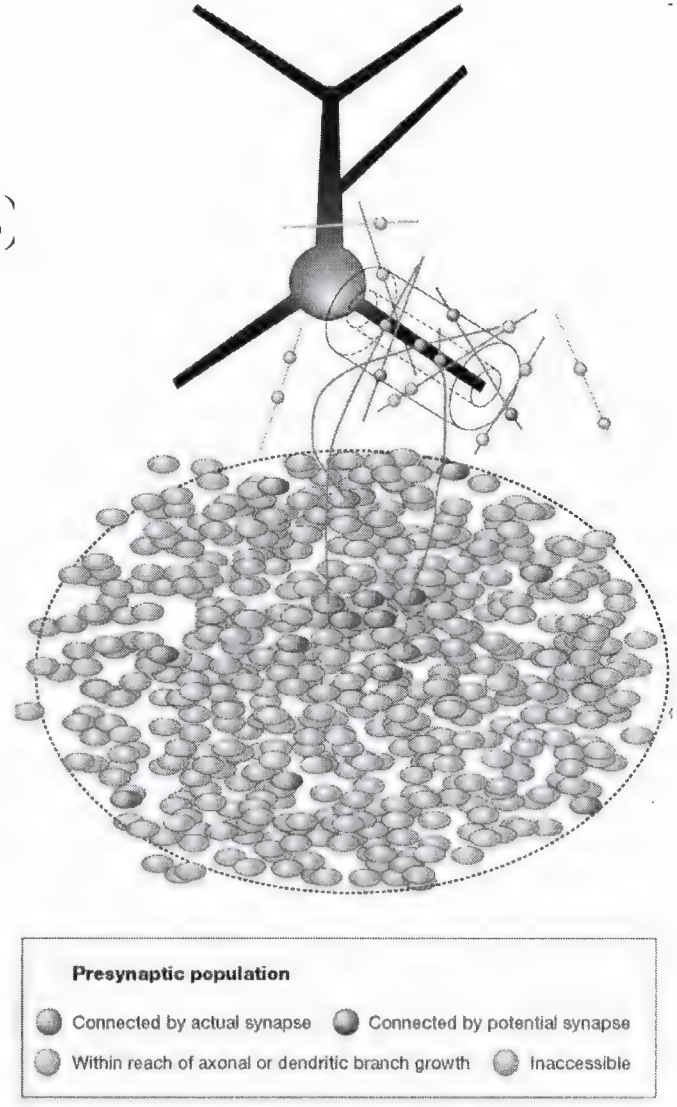

Chklovskii, Mel, Svoboda, 2004

Stability of synaptic structure in the adult brain

Boutons

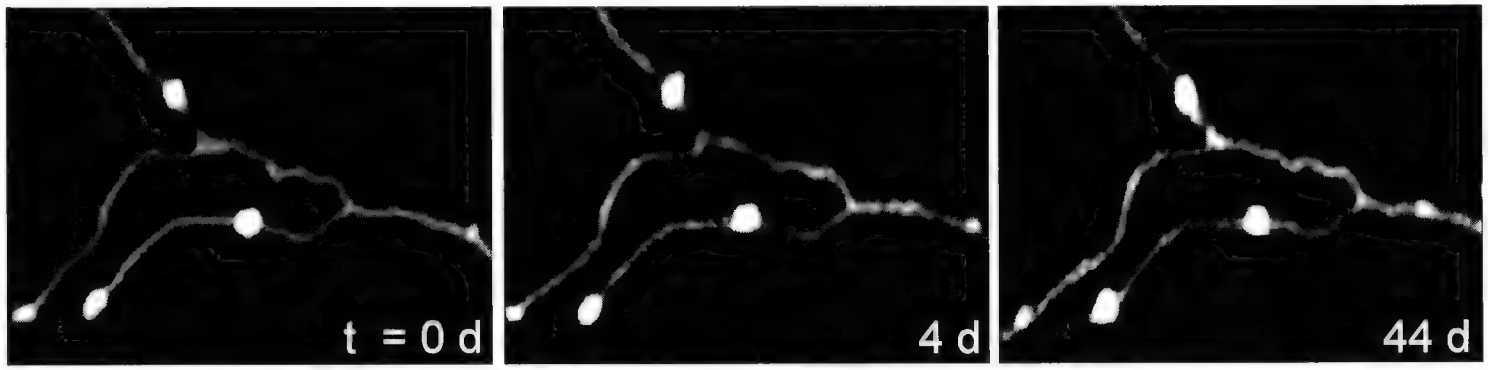

De Paola *et al. Neuron* 2006

Spines

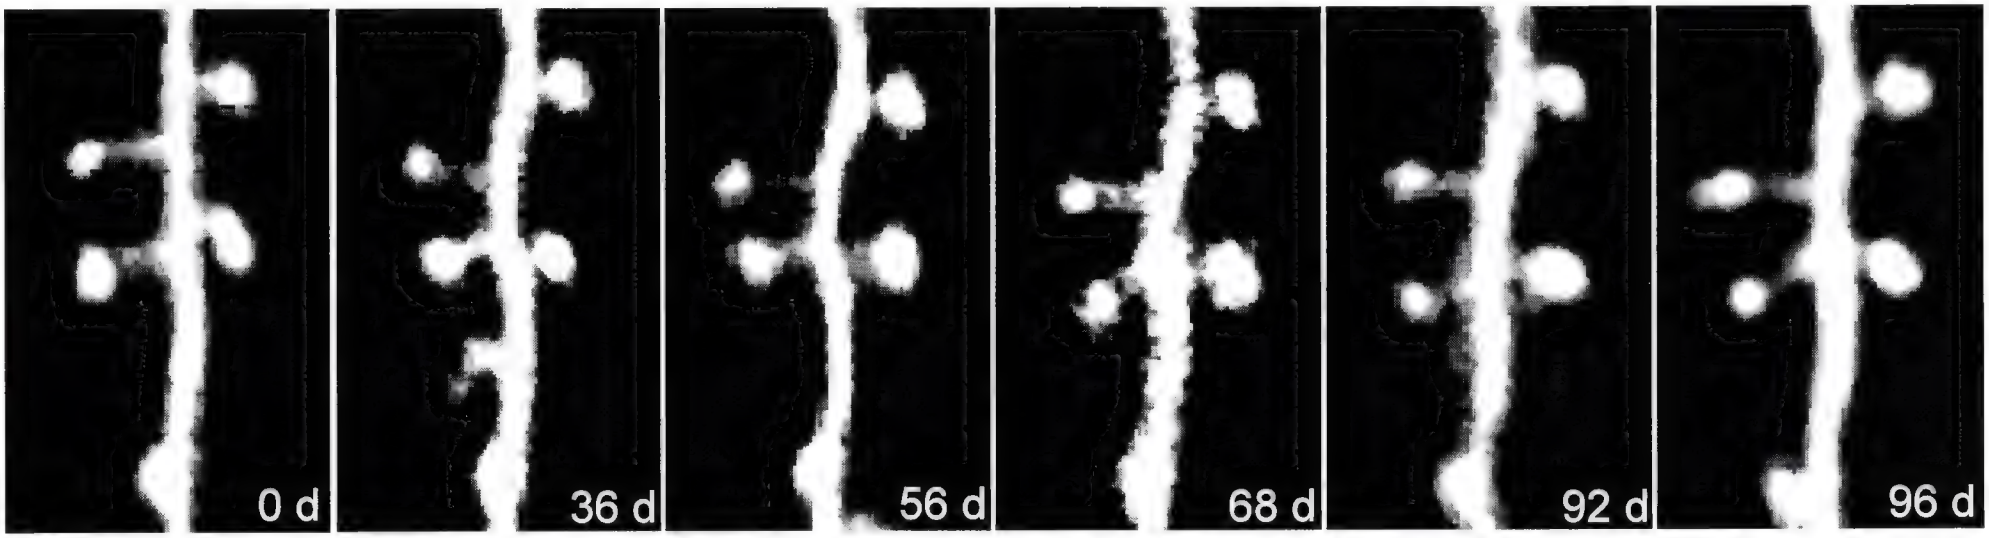

Trachtenberg *et al.*, *Nature* 2002, Grutzendler et al 2002; Hollmaat *et al.*, *Neuron* 2005

How can individual synapses maintain their size over months?

Learning induces long-term potentiation (LTP) of synaptic transmission in vivo

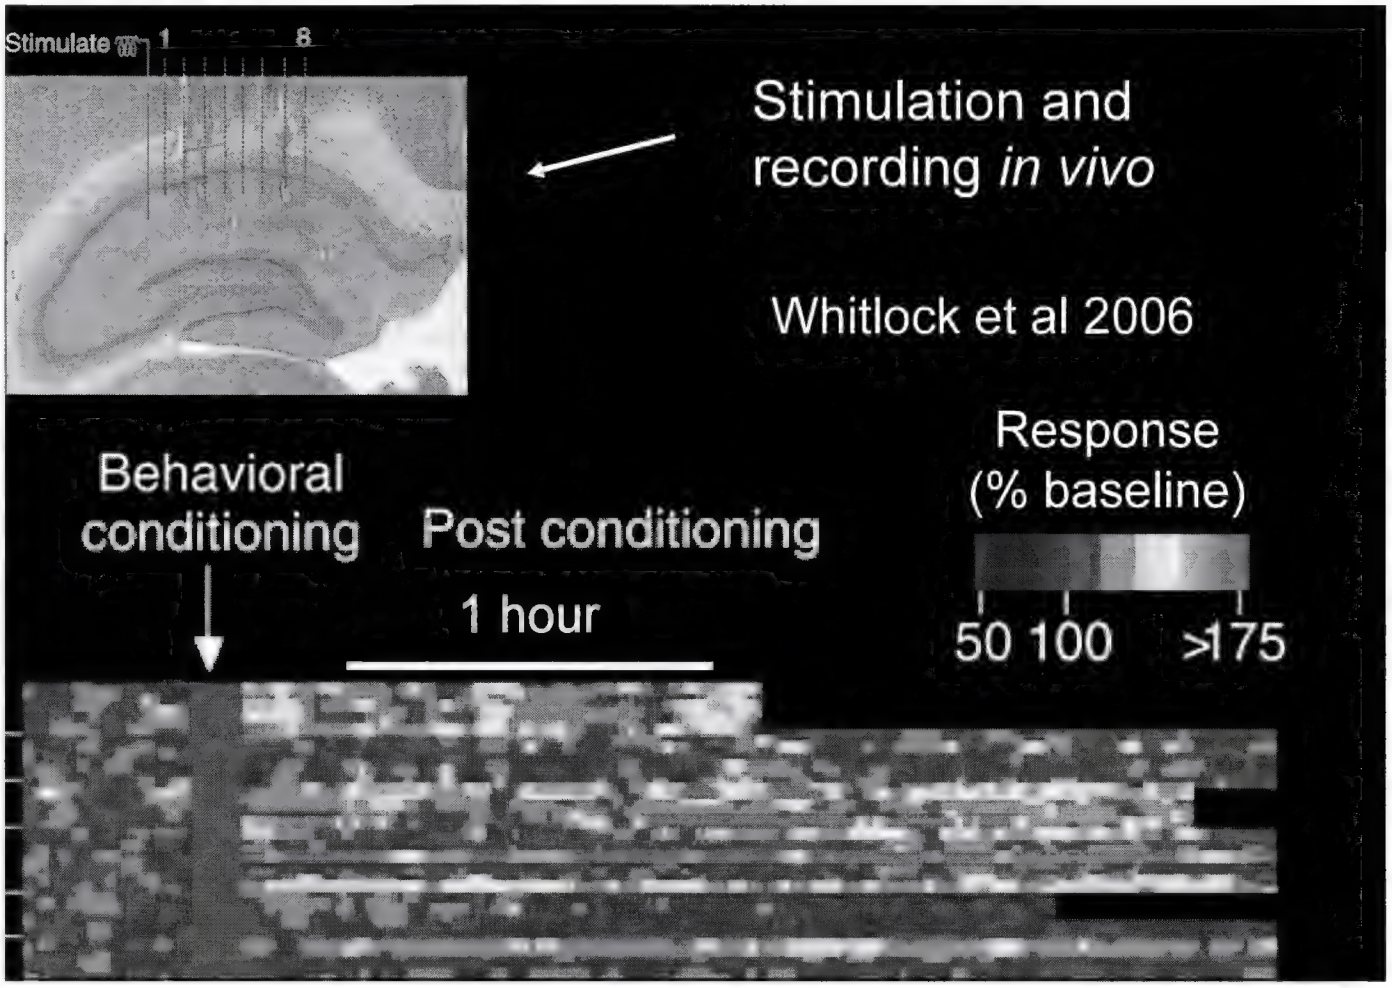

Lots of experiments to suggest that expression of synaptic plasticity is related to learning and memory

Long-term potentiation (LTP) of synaptic transmission

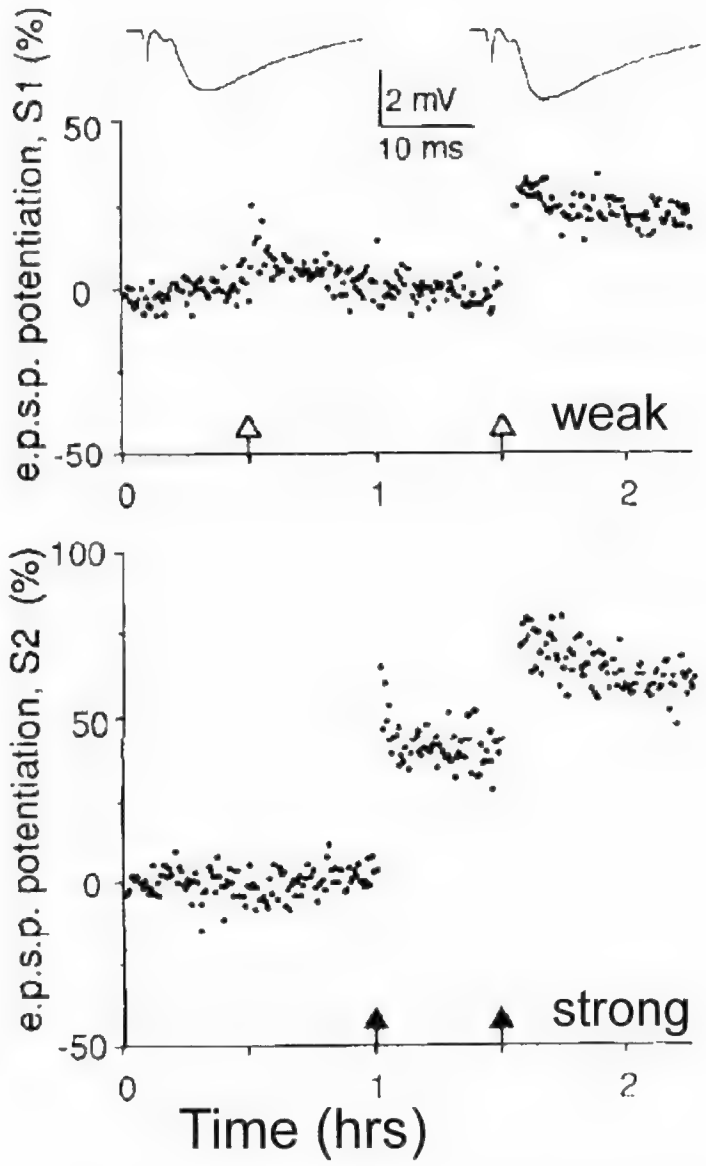

Bliss and Collingridge, 1993

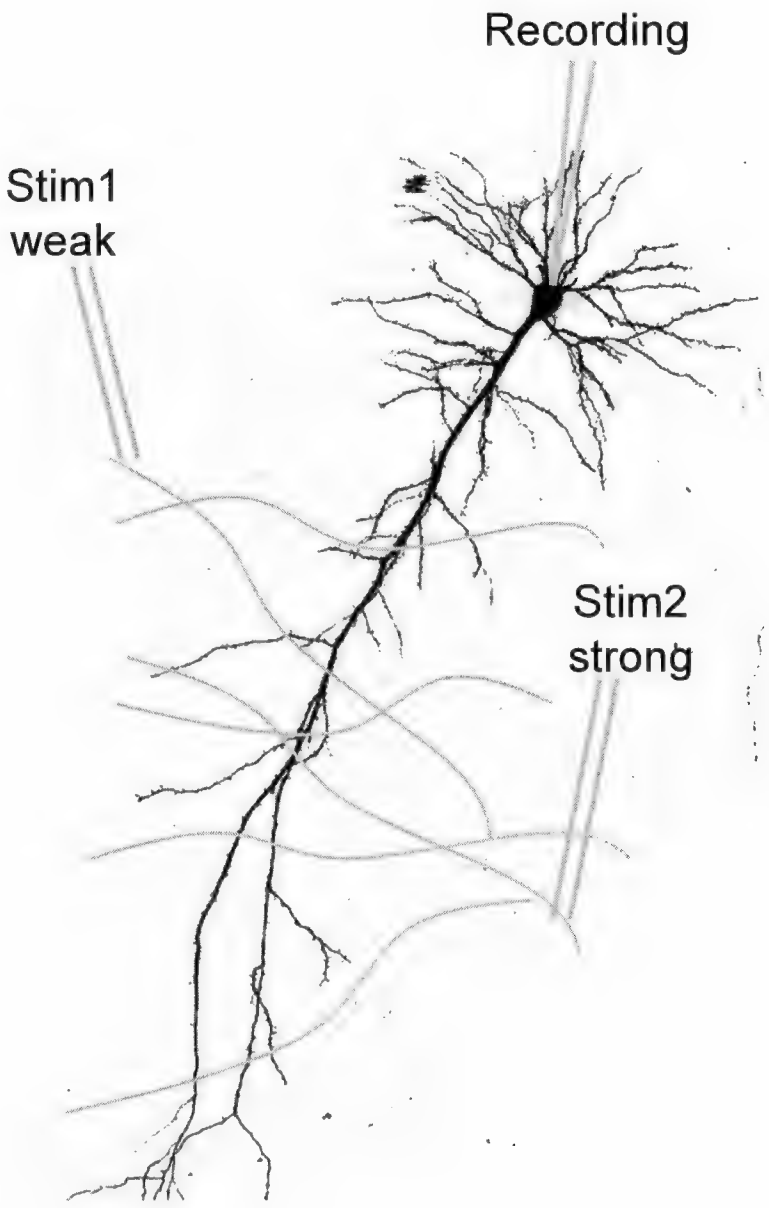

Induction protocols for LTP

- Tetanus

100 Hz, 1 second

- Theta burst

200 ms ISI, 5 pulses, 100 Hz etc

- Pairing

1 Hz, 1 minute, while depolarizing the postsynaptic cell

- STDP

Pair pre and postsynaptic action potentials –pre precedes post

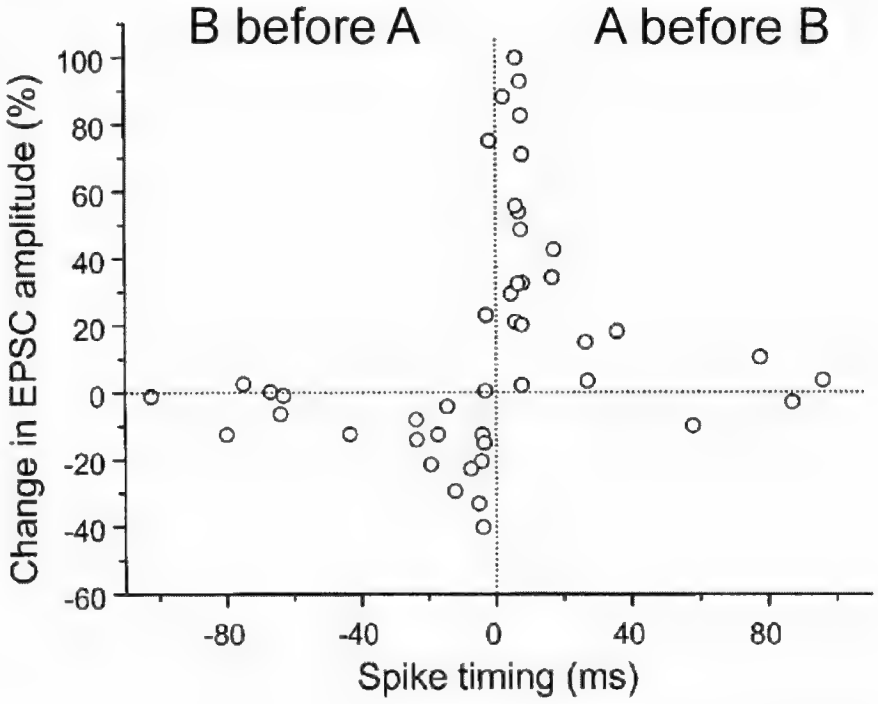

Bi and Poo 1998

Synaptic responses evoked by 2-photon glutamate uncaging

2-photon Ca^{2+} imaging &
2-photon glutamate uncaging

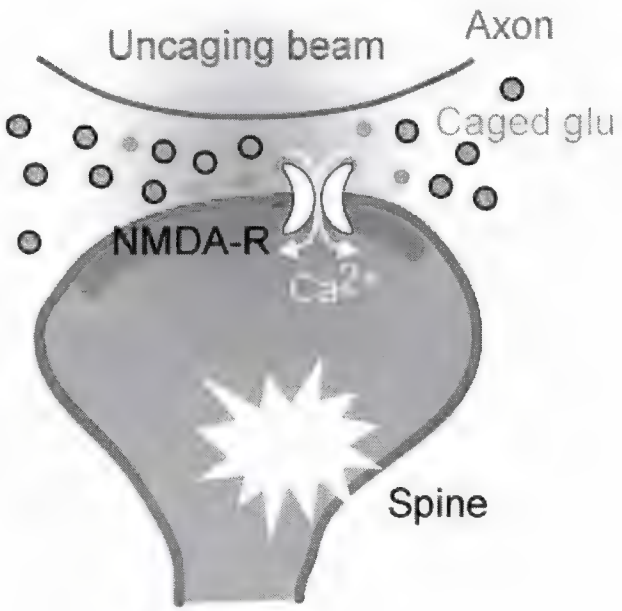

AMPA-R current

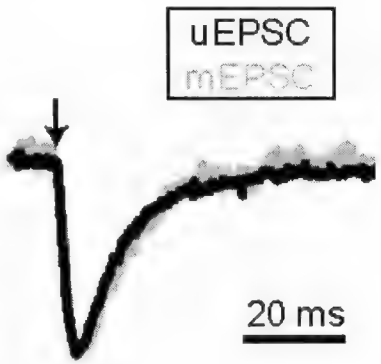

NMDA-R Calcium

Matsuzaki et al 2001
Carter and Sabatini, 2004
Sobczyk et al 2005

•Spine [Ca^{2+}] signals are synapse-specific

→ Is the induction of synaptic plasticity synapse-specific?

Long-Term Potentiation (LTP) is input specific

LTP protocol (0.5 Hz, 30 pulses)

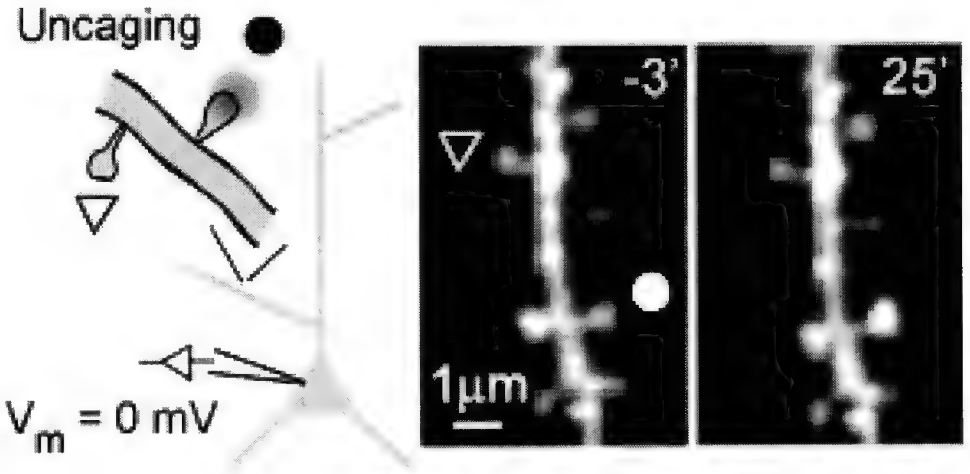

Long-term potentiation (LTP)

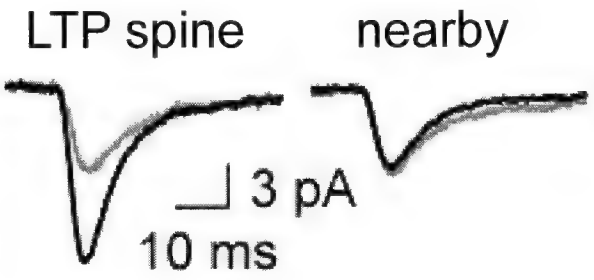

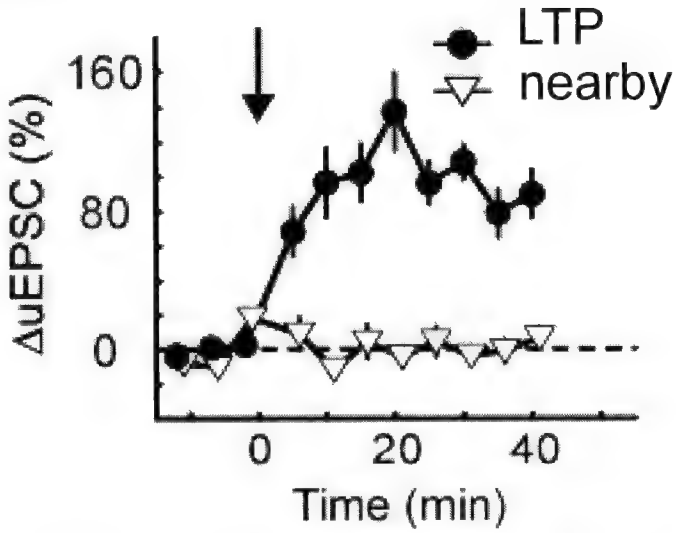

Matsuzaki et al 2004; Harvey and Svoboda 2007

Spine enlargement is a structural correlate of LTP

LTP protocol (0.5 Hz, 30 pulses)

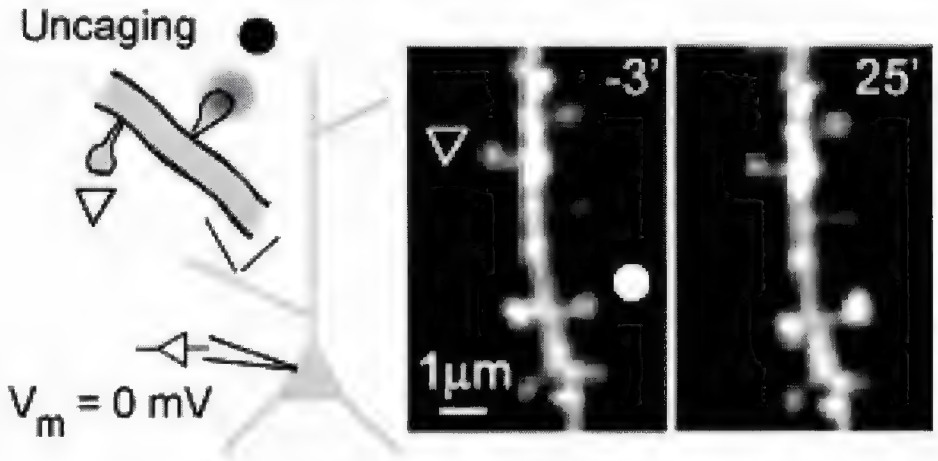

Spine enlargement

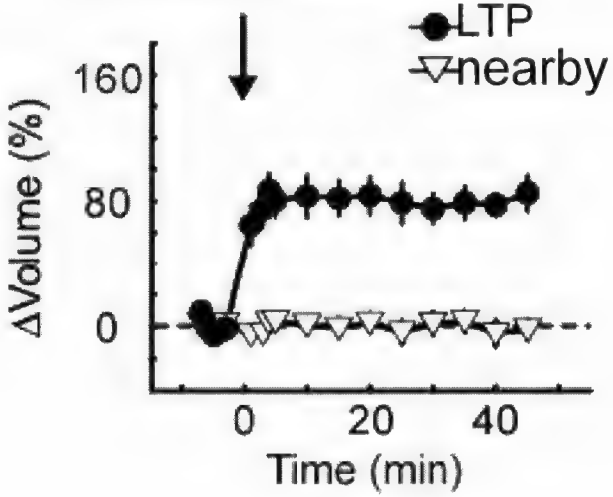

Spine enlargement vs LTP

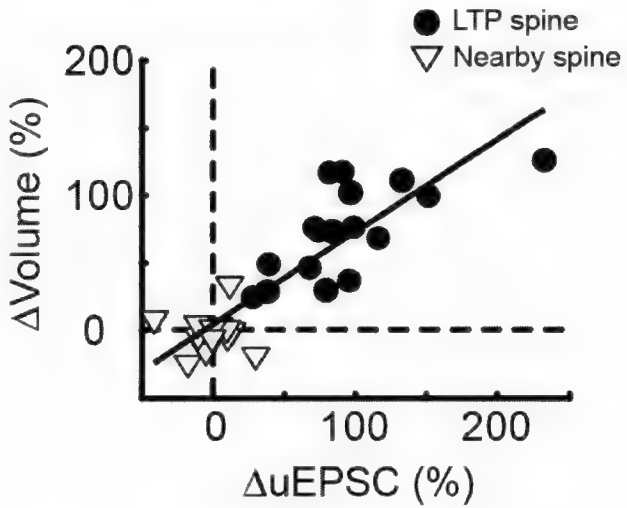

Matsuzaki et al 2004; Harvey and Svoboda, 2007

Synapse-specific signaling and plasticity

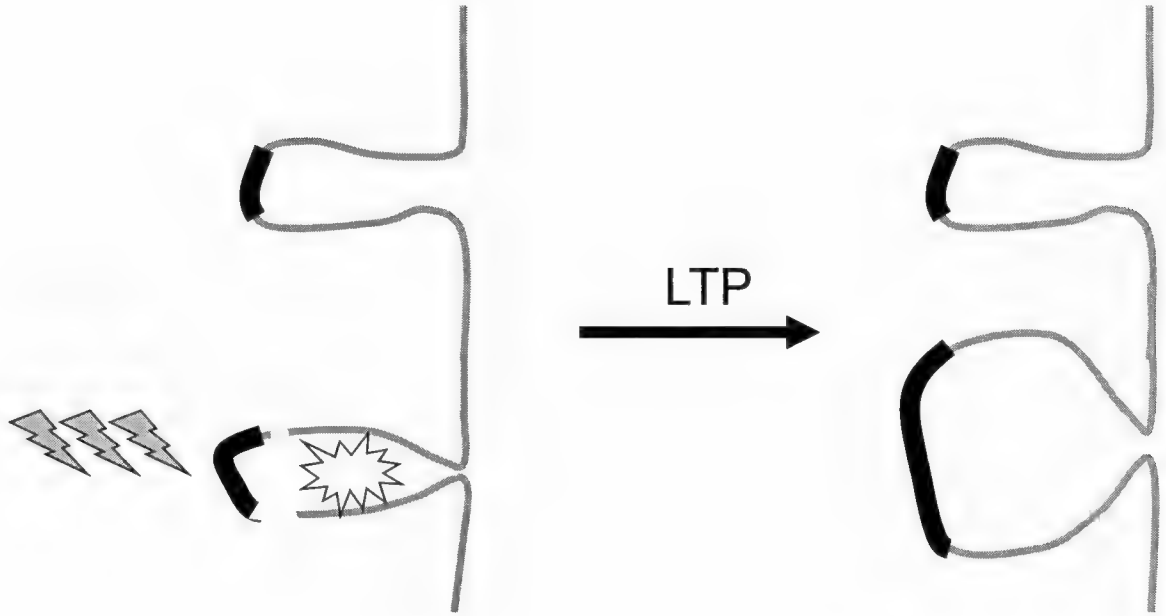

- Synaptic transmission is point-to-point
- Spine [Ca^{2+}] signals are synapse-specific
- LTP and spine enlargement are synapse-specific

→ Synapses can function as independent units for information storage

Calcium-dependent signal transduction in spines

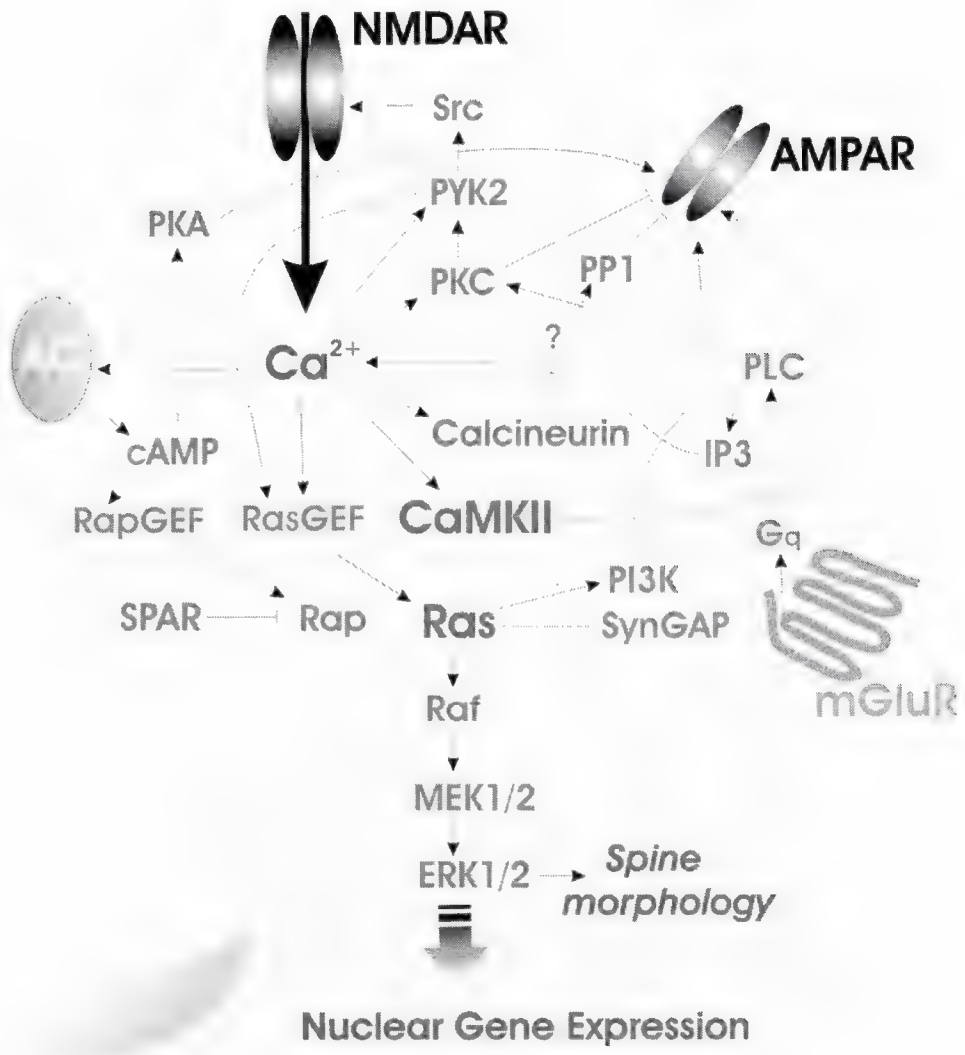

Sheng M & Kim MJ, 2002

The Ras pathway

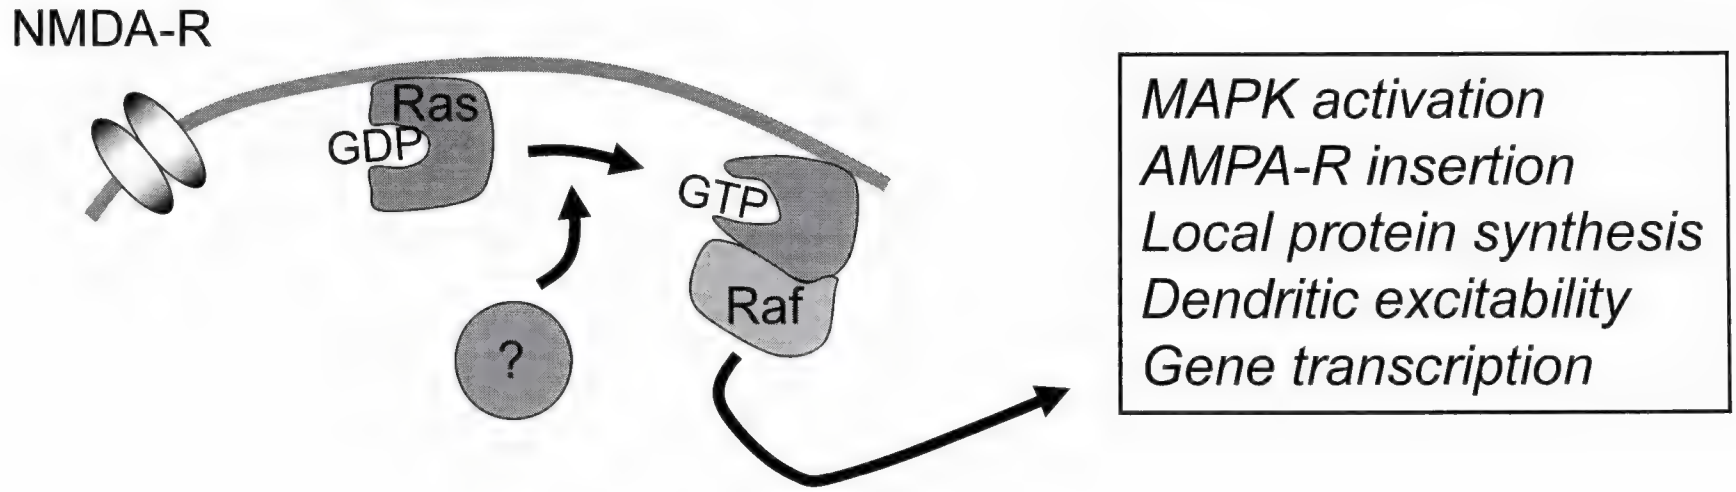

- Downstream of NMDA-R activation and Ca^{2+}
- Ras & GAPs, GEFs concentrated in spines
- Critical for long-term potentiation
- Critical for spine-enlargement
- Ras pathway mutants are implicated in mental retardation
- Amenable by FRET analysis

Ras activity measured using fluorescence lifetime imaging (FLIM)

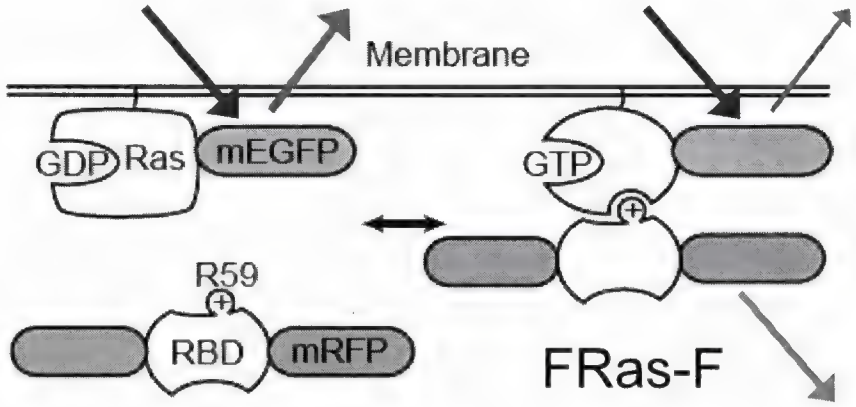

- Only track donor fluorescence
- Quantitative (binding fraction)
- Real-time
- Compatible with 2-photon microscopy
- Applicable to scattering tissue

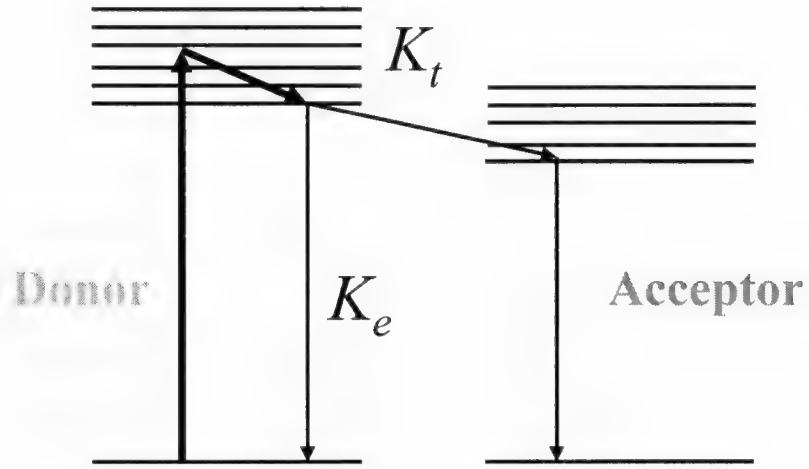

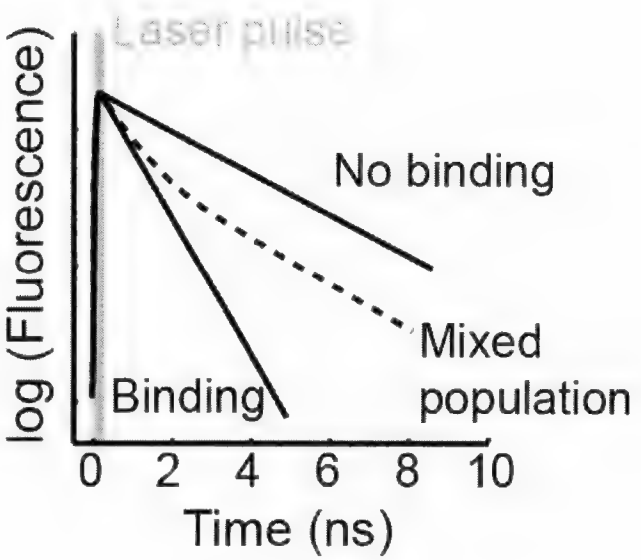

Ras sensor (FRas-F) reports Ras activation

Response of Ras sensor

Imaging

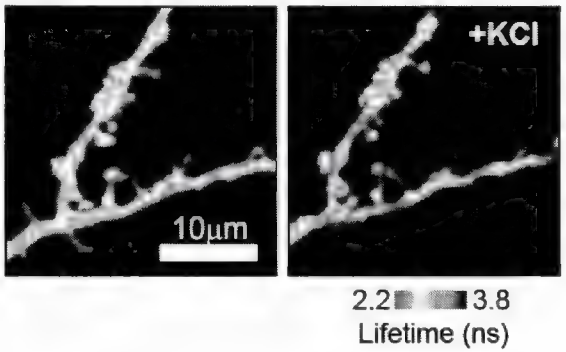

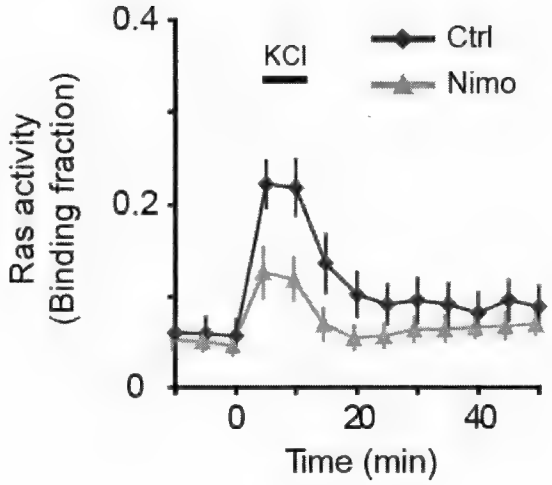

Response of endogenous Ras

Pull-down assay

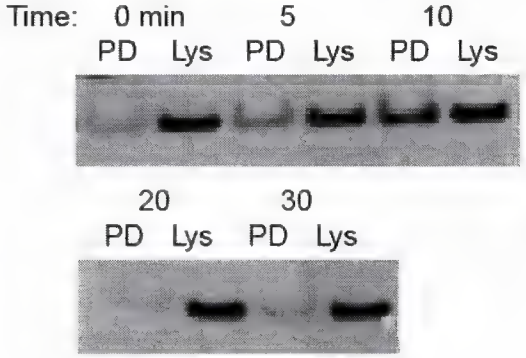

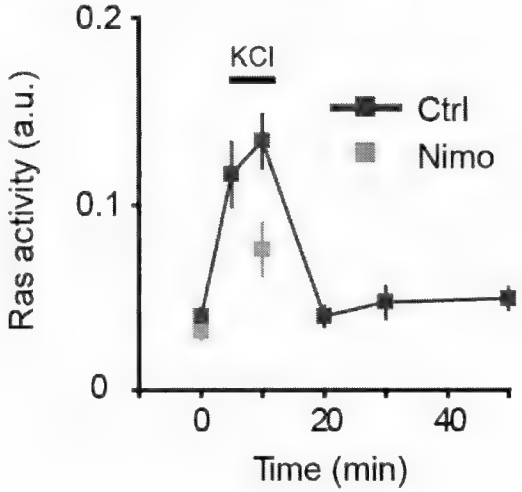

Tune sensor for:

- FLIM
- Signal level
- Biology

Yasuda et al 2006

Ras activity in single spines during plasticity induction

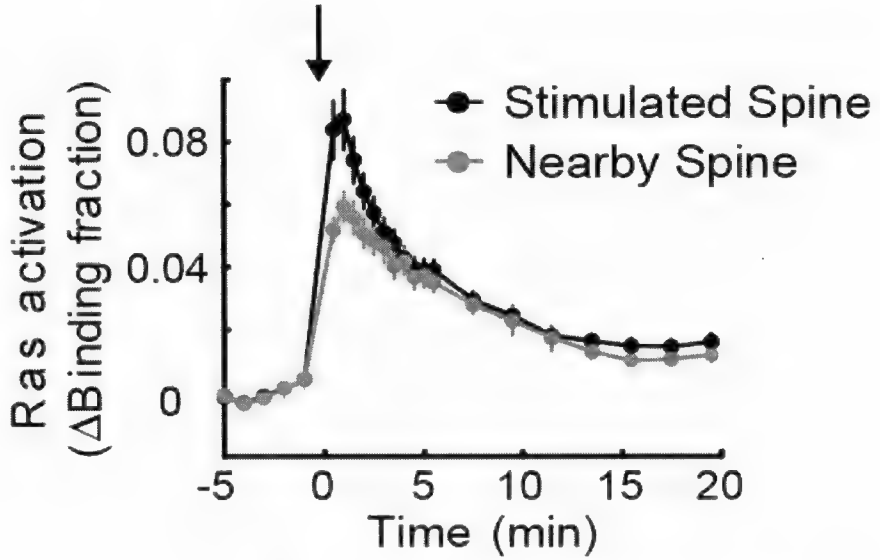

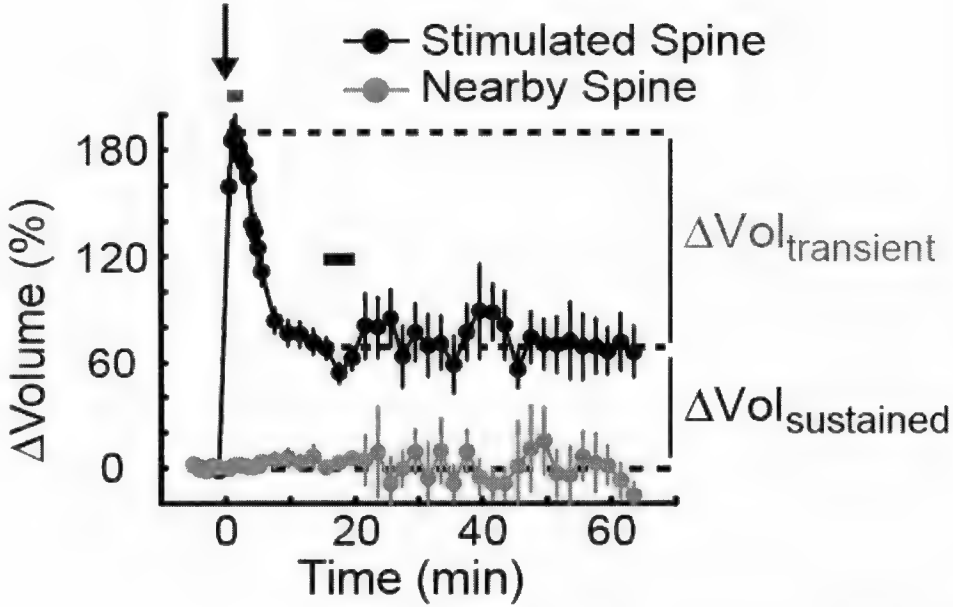

2pFLIM image (1/min)

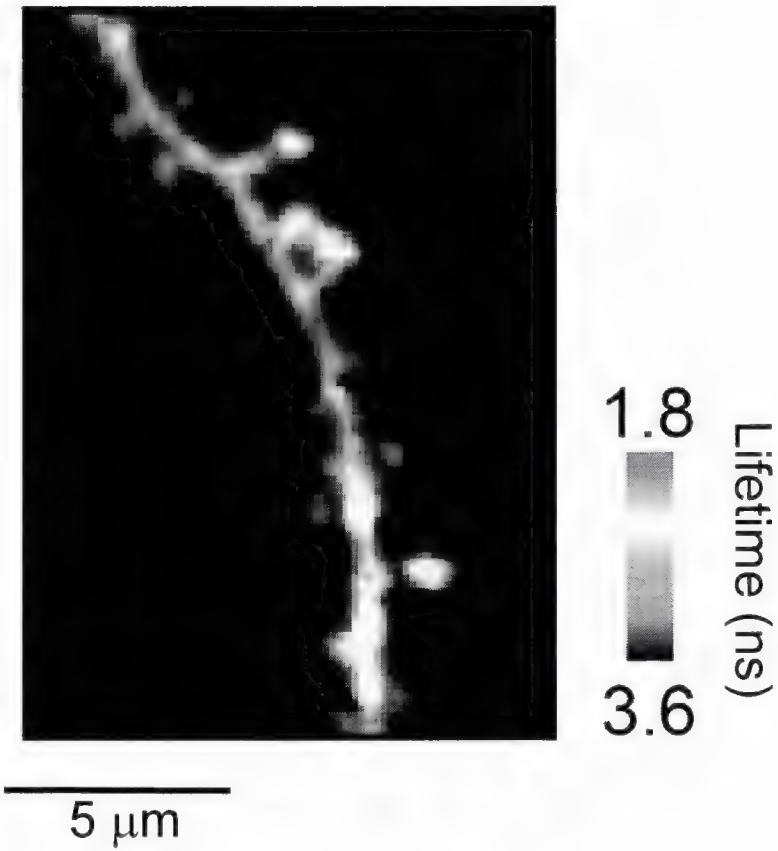

Harvey, Yasuda *et al* 2008

Spatial spread of Ras activity

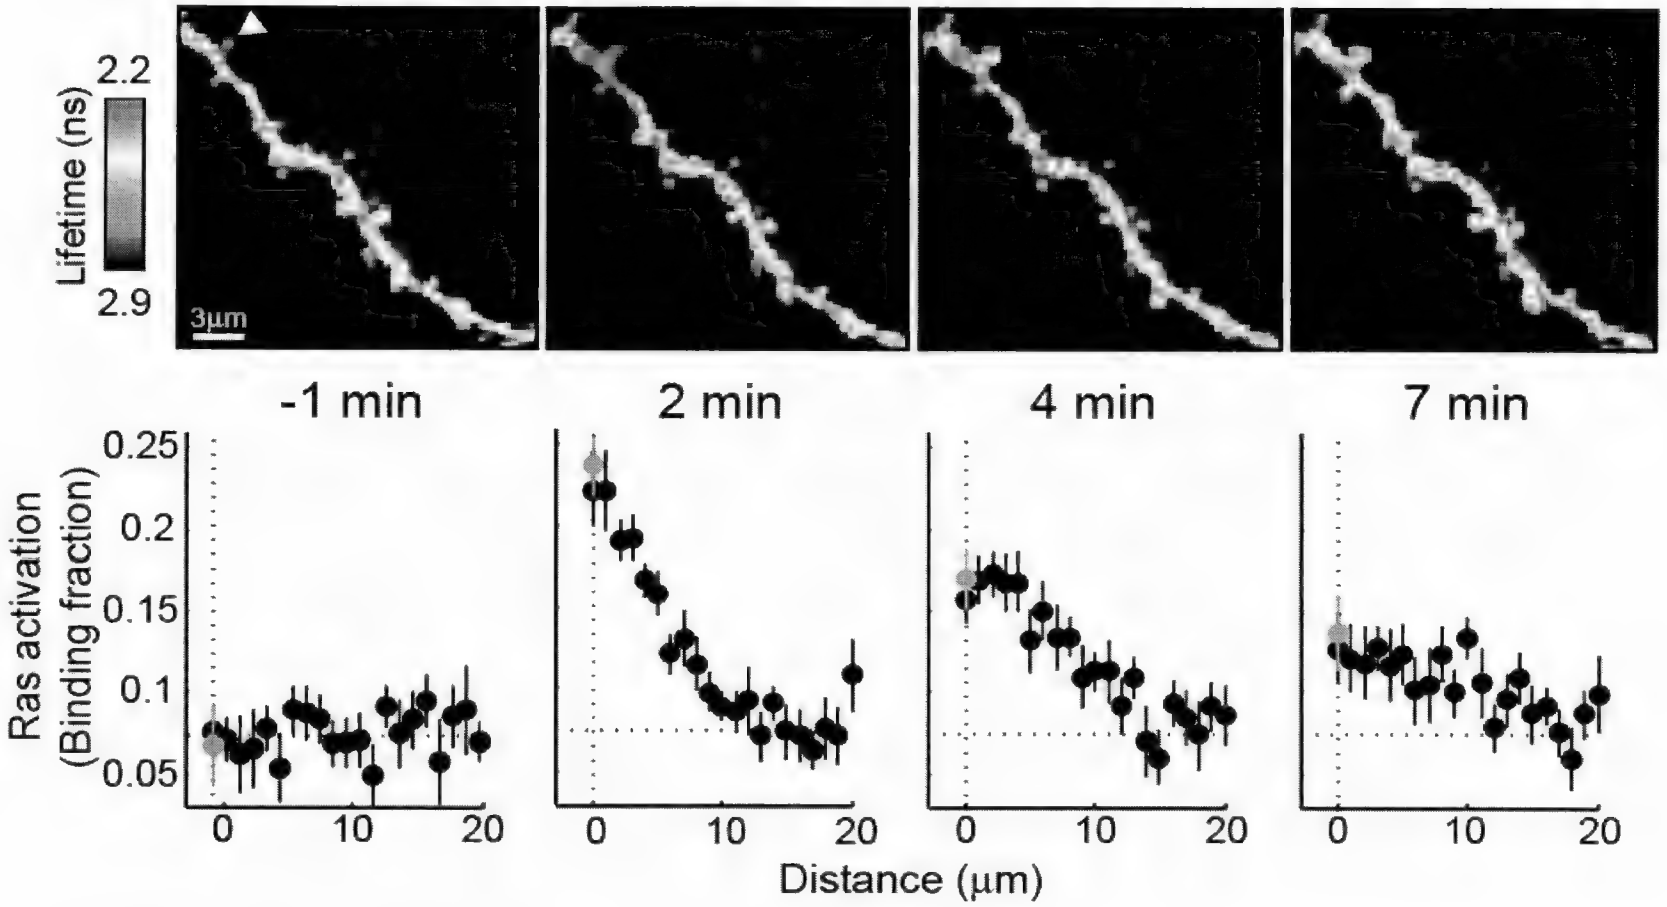

- Active Ras spreads ~ 10 μm by diffusion ($L \sim (D_{Ras}\ \tau_{inactivation})^{1/2}$)
- … corresponding to multiple synapses

→ Ras signaling is not synapse-specific

CaMKII activity is spine-specific

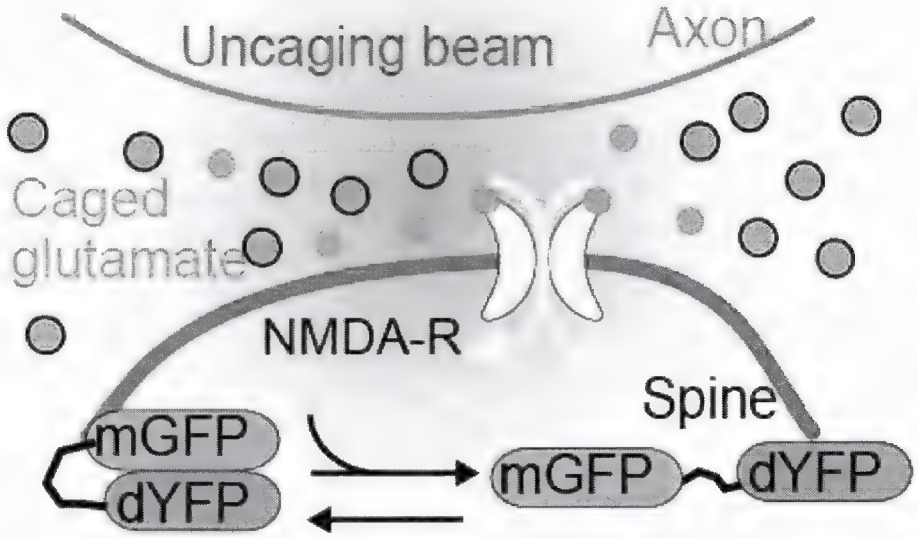

Lee et al 2009; Yasuda lab

2pFLIM image

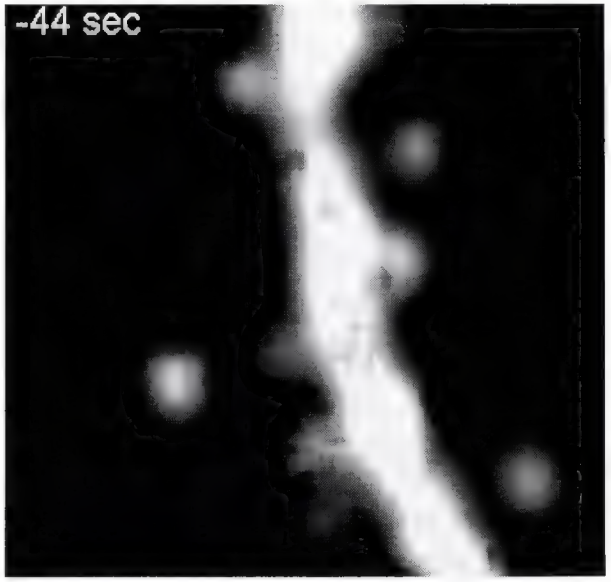

•Localized CaMKII activity likely underlies synapse-specific LTP

Synapse - specificity in Ca^{2+}-dependent signaling

	Ca^{2+}	**CaMKII**	**Ras**
Extrusion / inactivation time (τ_{off})	10 ms	30 sec	5 min
Retention time, spine (τ_r)	>100 ms	10 min	1-10 sec
(τ_{off}) / (τ_r)	<< 1	<< 1	>> 1
Synapse-specific	**Yes**	**Yes**	**No**

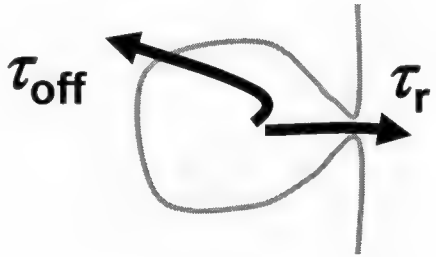

Sabatini et al. 2002
Harvey, Yasuda *et al* 2008
Lee et al 2009

Ras signaling is not synapse-specific

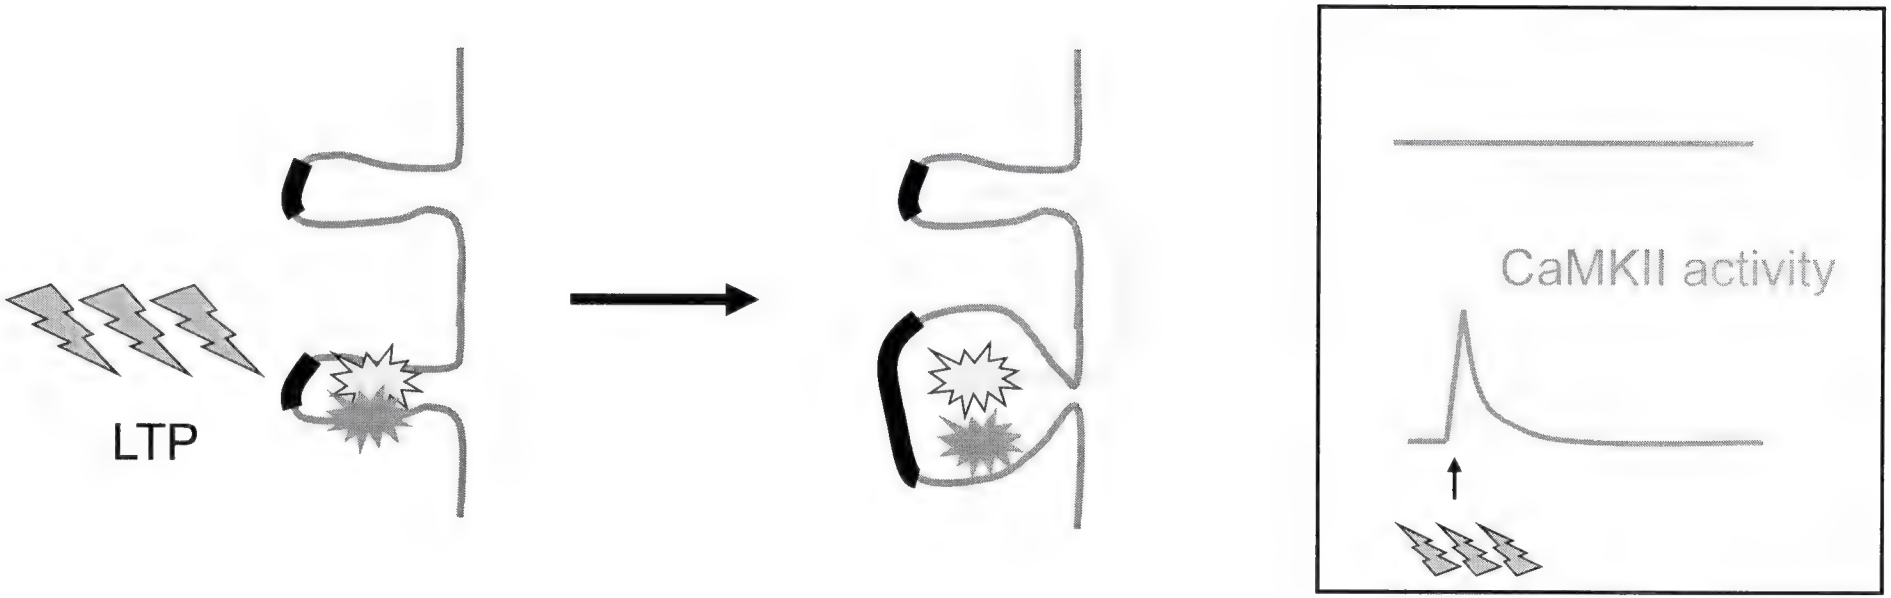

- Localized CaMKII activity likely underlies synapse-specific LTP

→ **What are the consequences of the spatial spread of Ras?**

Crosstalk in synaptic plasticity between nearby synapses

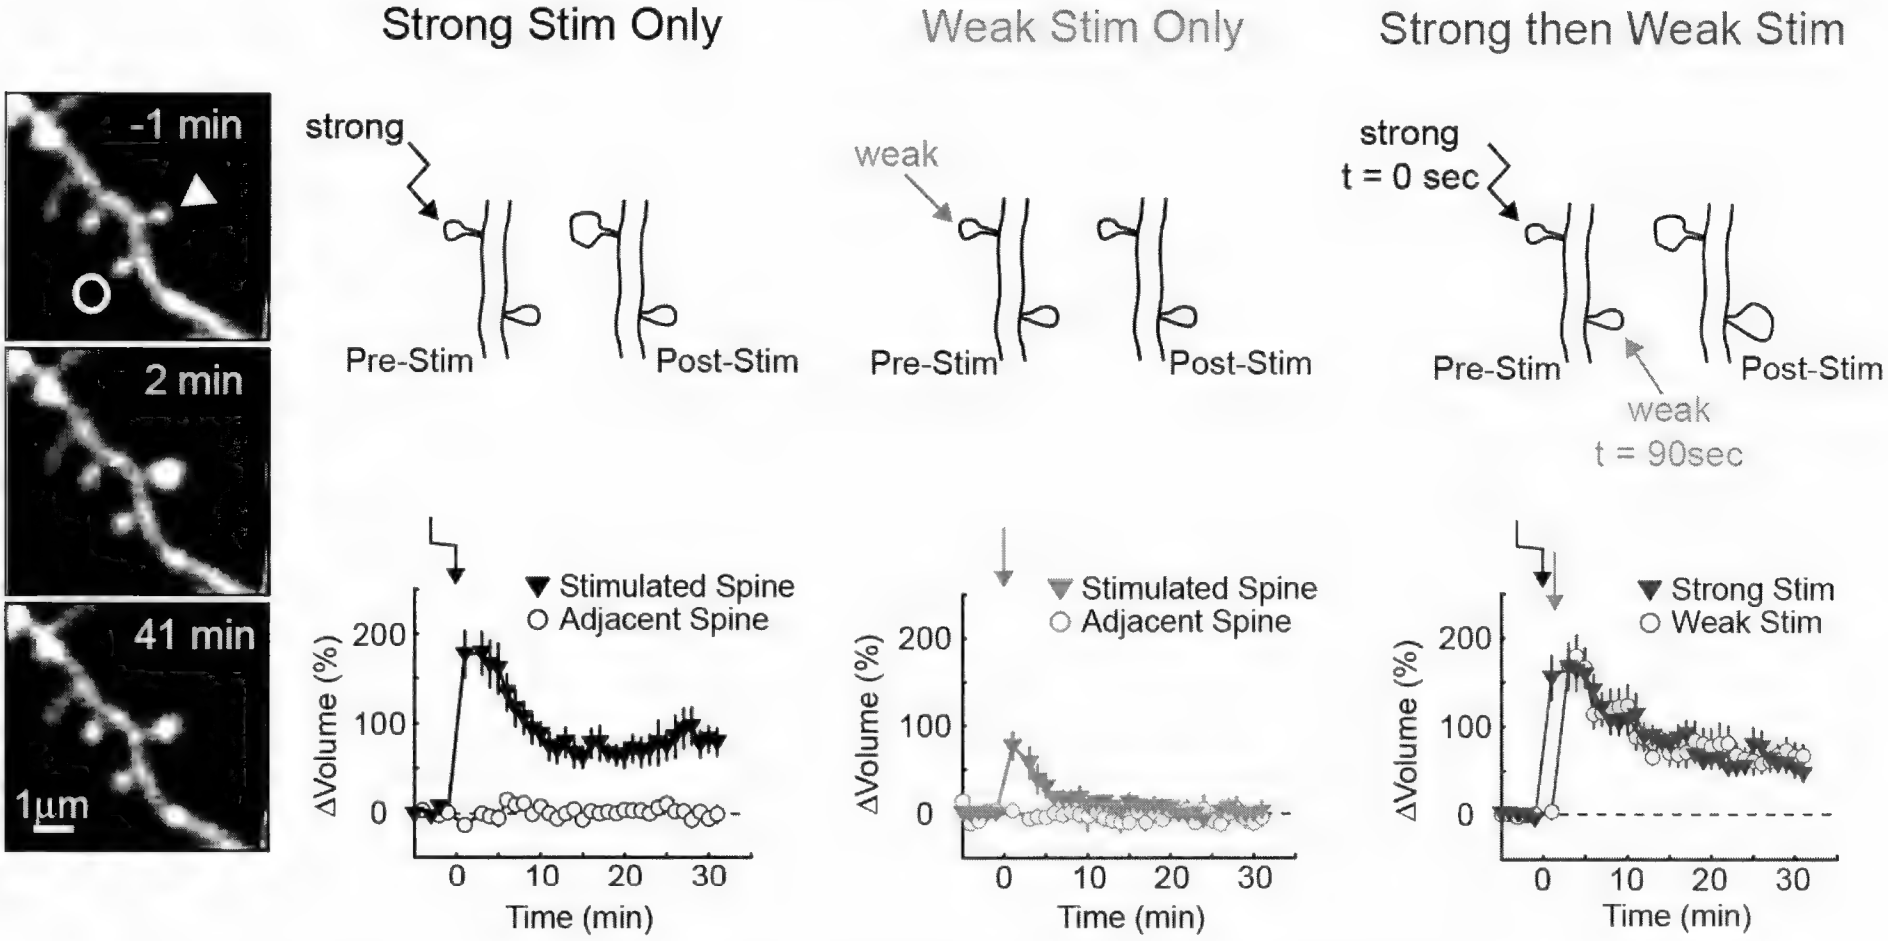

- Induction of LTP and spine enlargement is input-specific
- Induction of plasticity at one synapse facilitates induction at neighboring synapses (crosstalk)

Harvey and Svoboda, 2007

Crosstalk and Ras activity

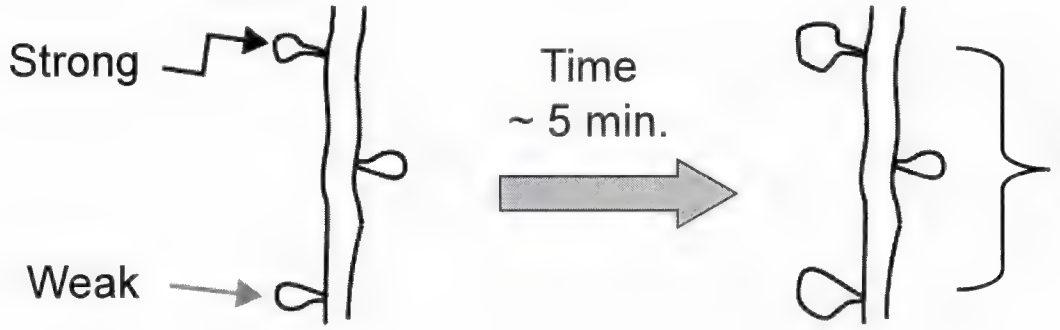

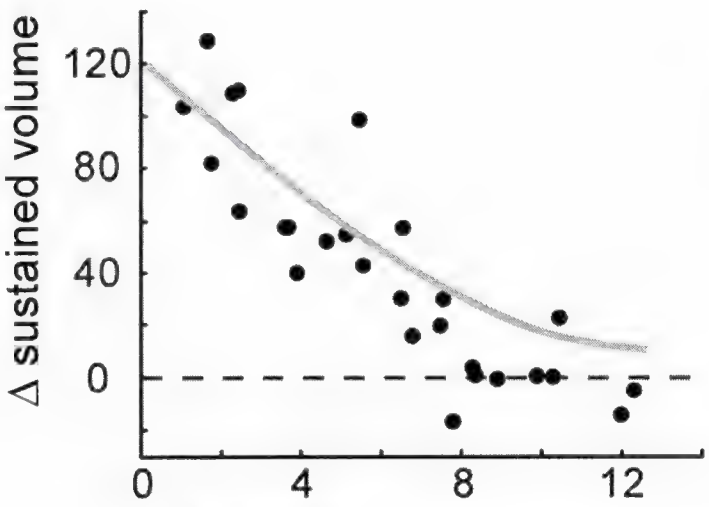

Distance from strongly stimulated spine (μm)

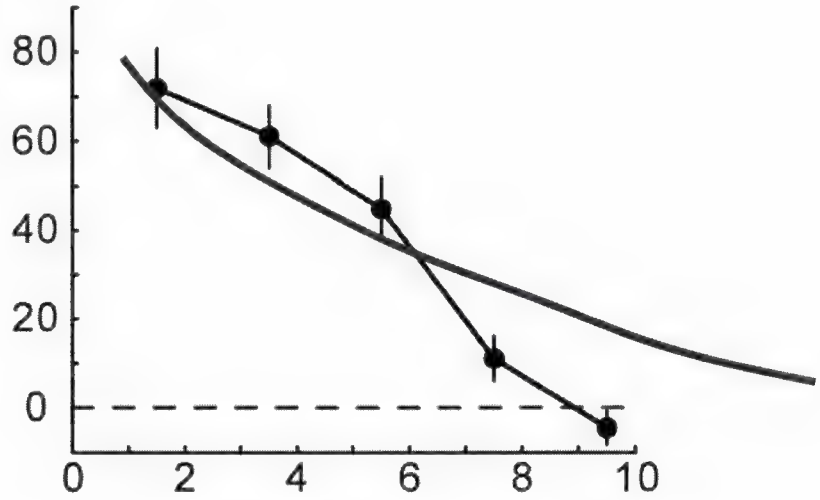

Time between strong and weak stimuli (minutes)

The spread of Ras activity is necessary for crosstalk

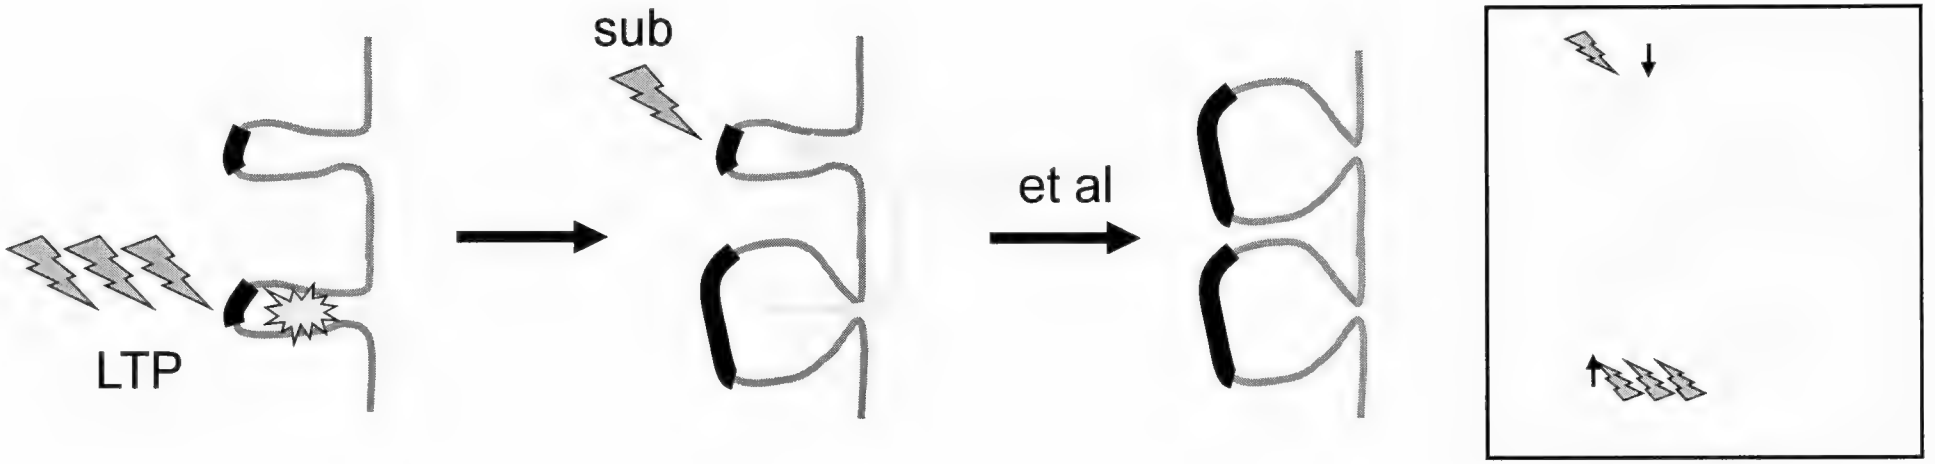

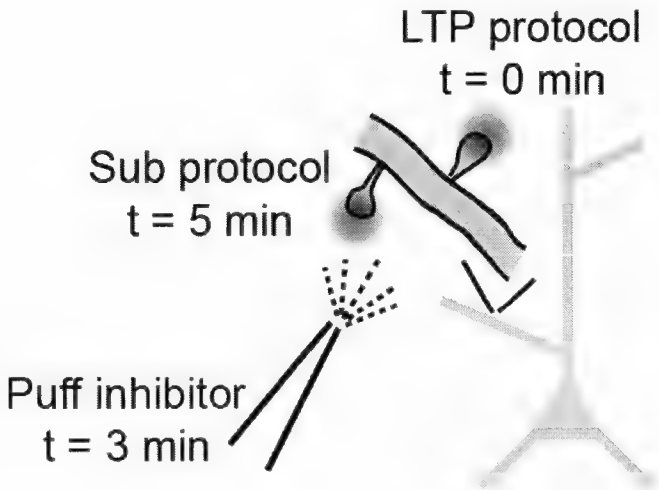

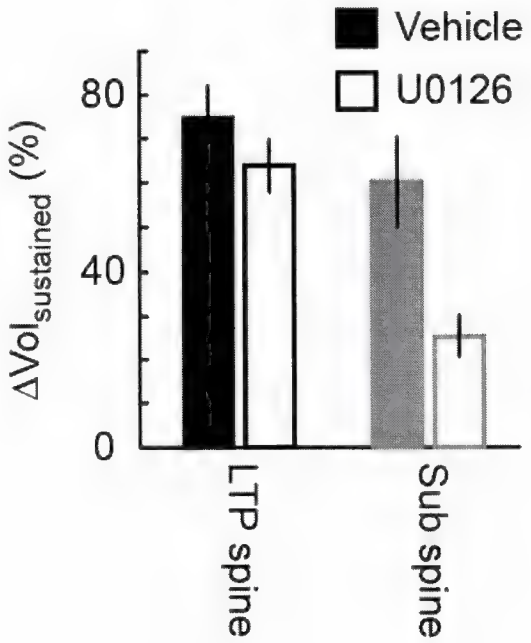

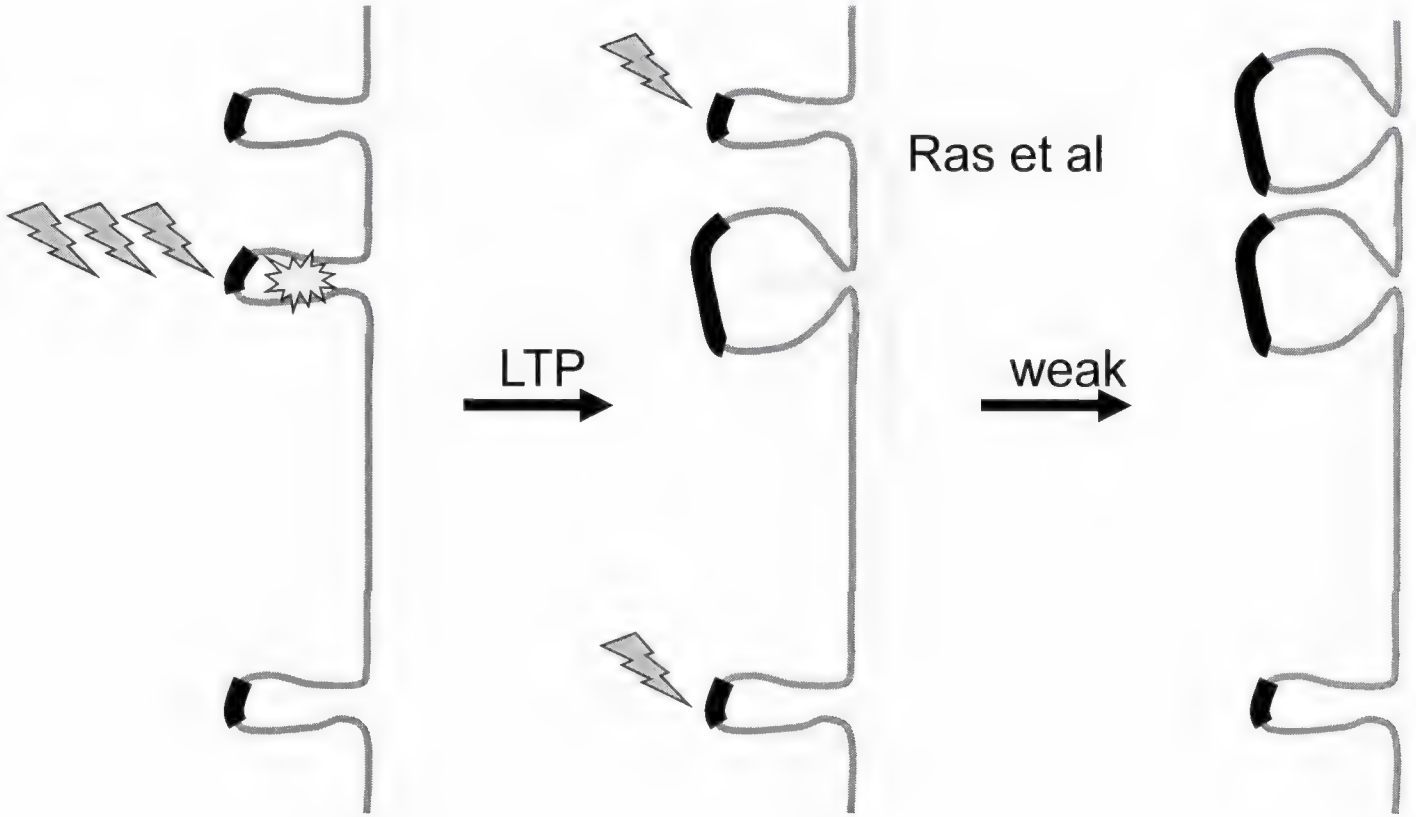

•Plasticity at one synapse facilitates induction at neighboring synapses

•Synapses are coupled by diffusible activated factors, including Ras

•Crosstalk over ~ 10 micrometers (10-20 synapses) and up to 5 minutes! Behaviorally relevant timescales …

Speculation ….

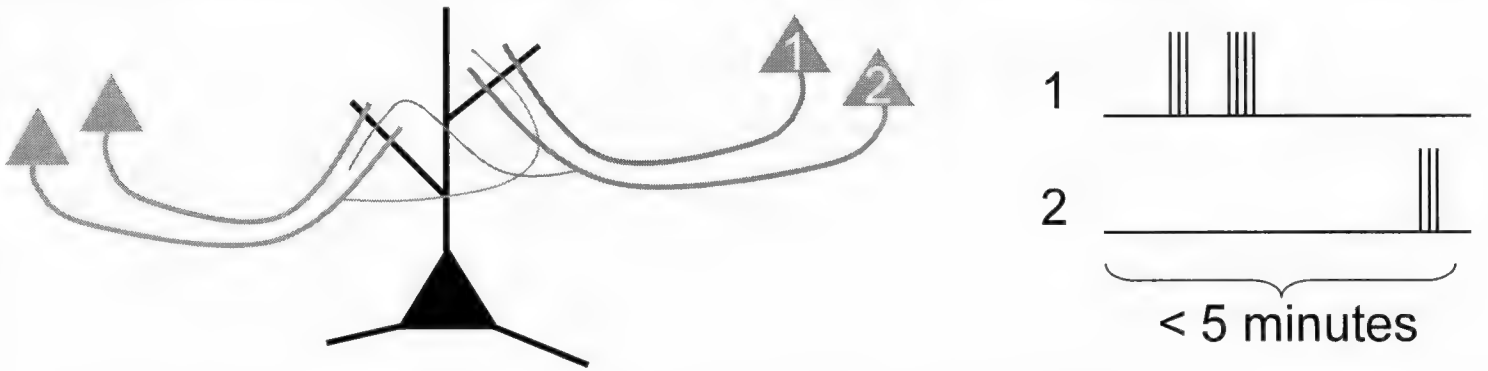

Inputs that are coactive during brief epochs become focused within the dendritic arbors of their target neurons

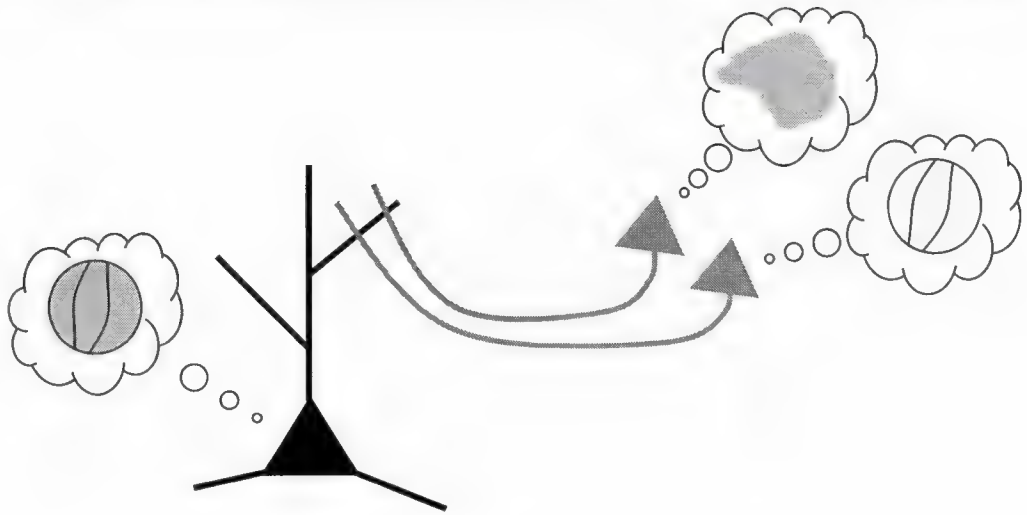

Mechanism for spatial clustering of semantically related inputs … ?

Barrel cortex: Inputs to pyramidal cells are arrayed in orderly, non-over-lapping dendritic domains

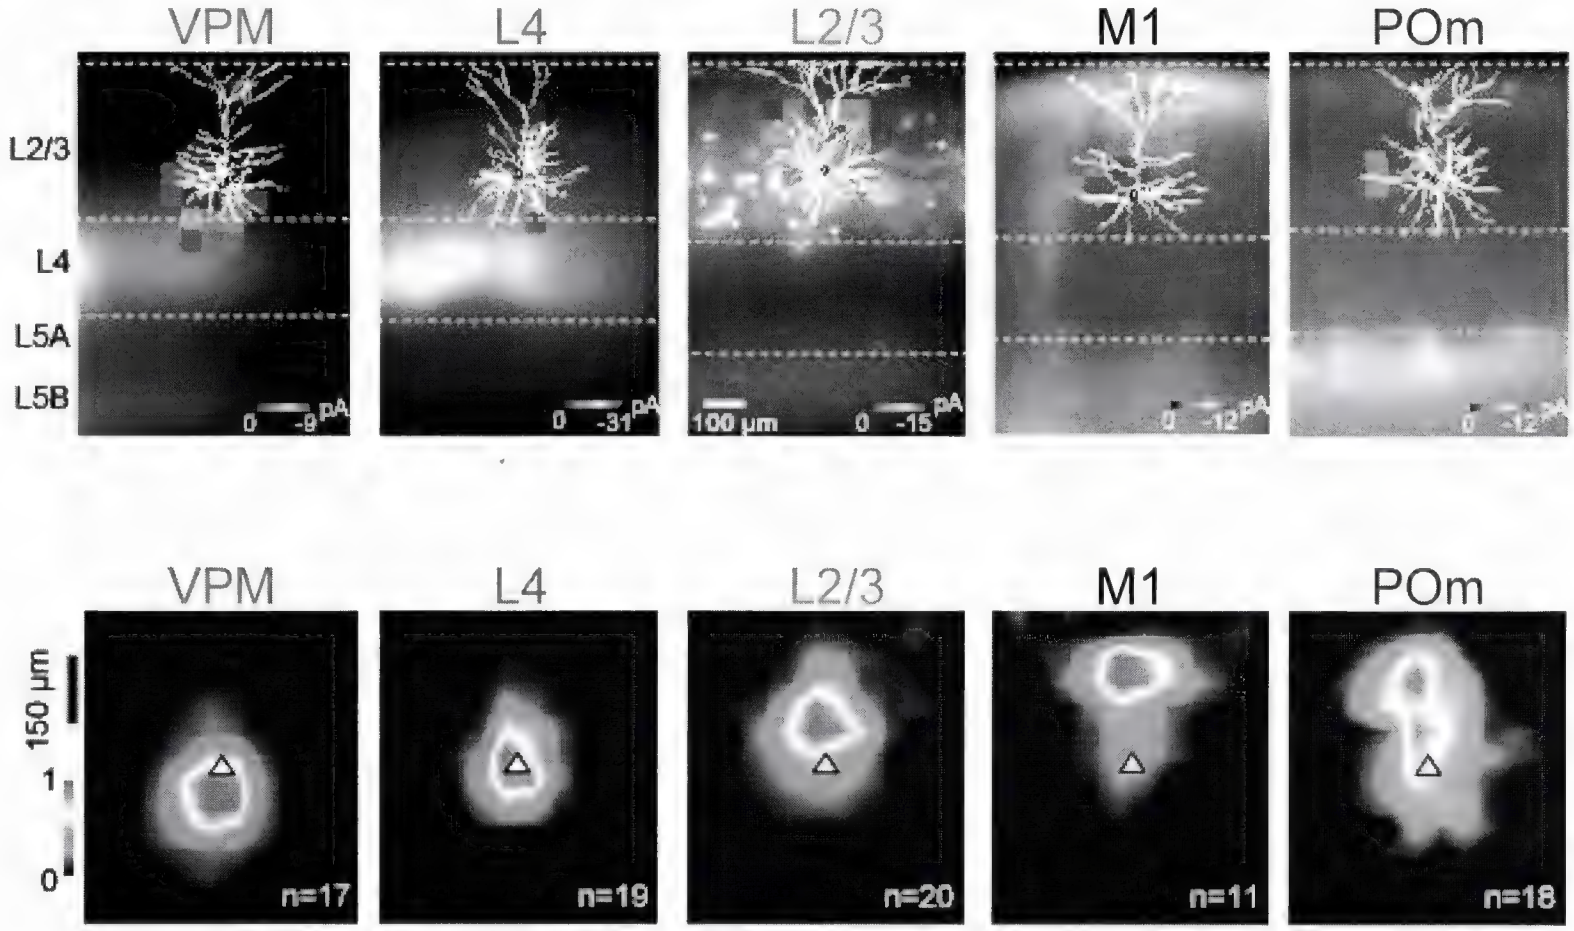

Leopoldo Petreanu

The circuits underlying somatosensation

- Ordered specific inputs within dendritic arbors
- Sensorimotor integration in single neurons?

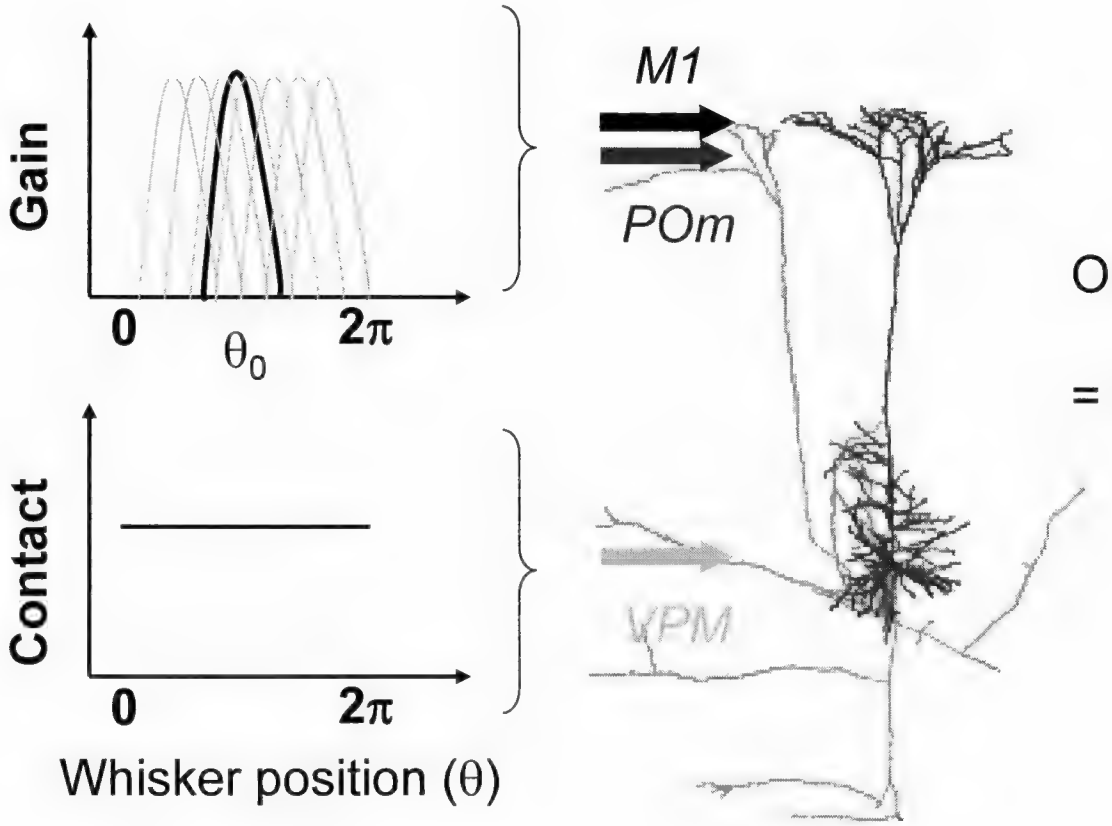

Output ==**Object location** at θ_0

= Gain (θ) * Contact

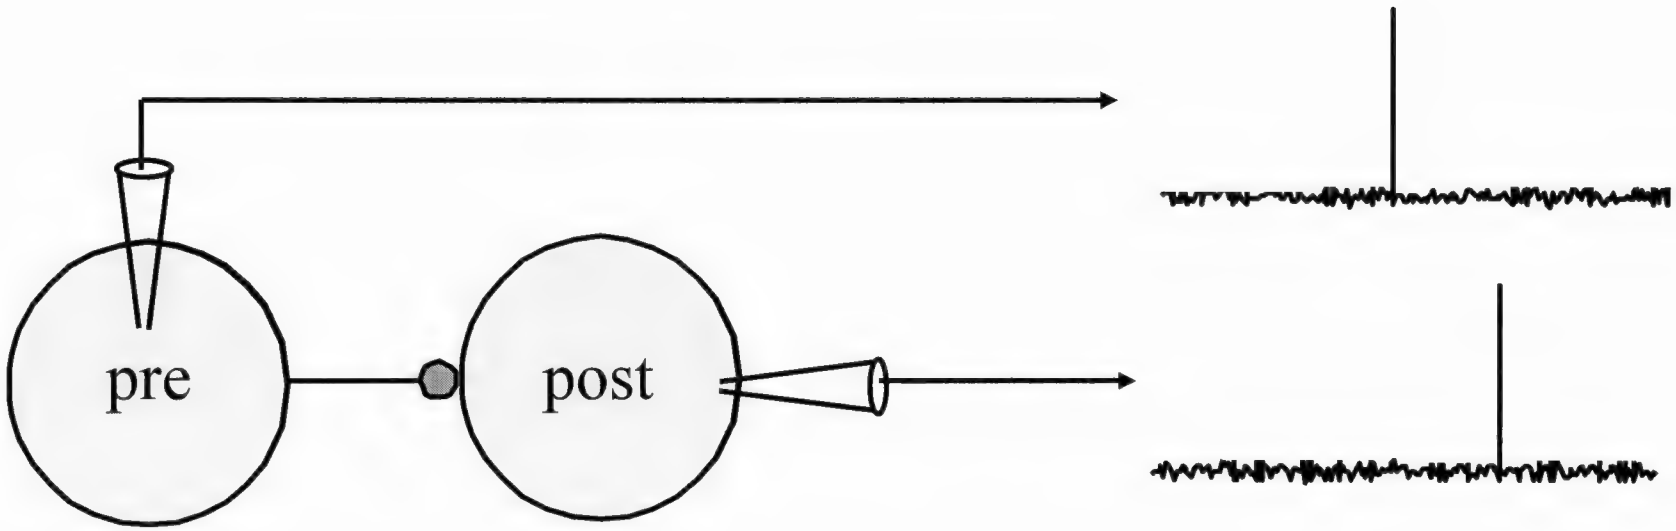
pre
post

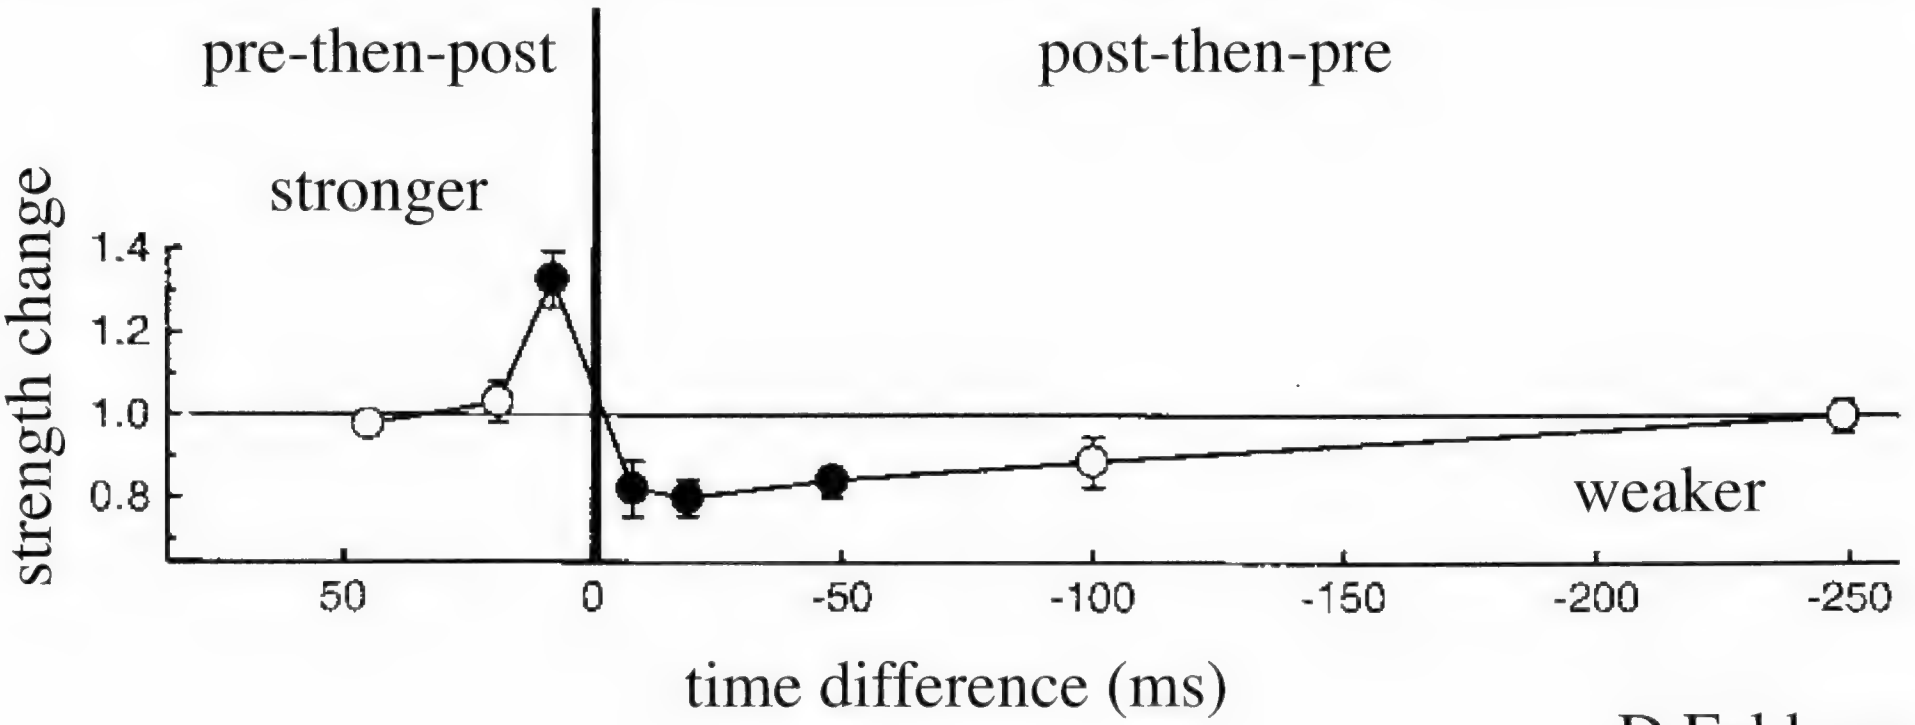
pre-then-post
post-then-pre
stronger
weaker
strength change
1.4
1.2
1.0
0.8
50
0
-50
-100
-150
-200
-250
time difference (ms)

D Feldman

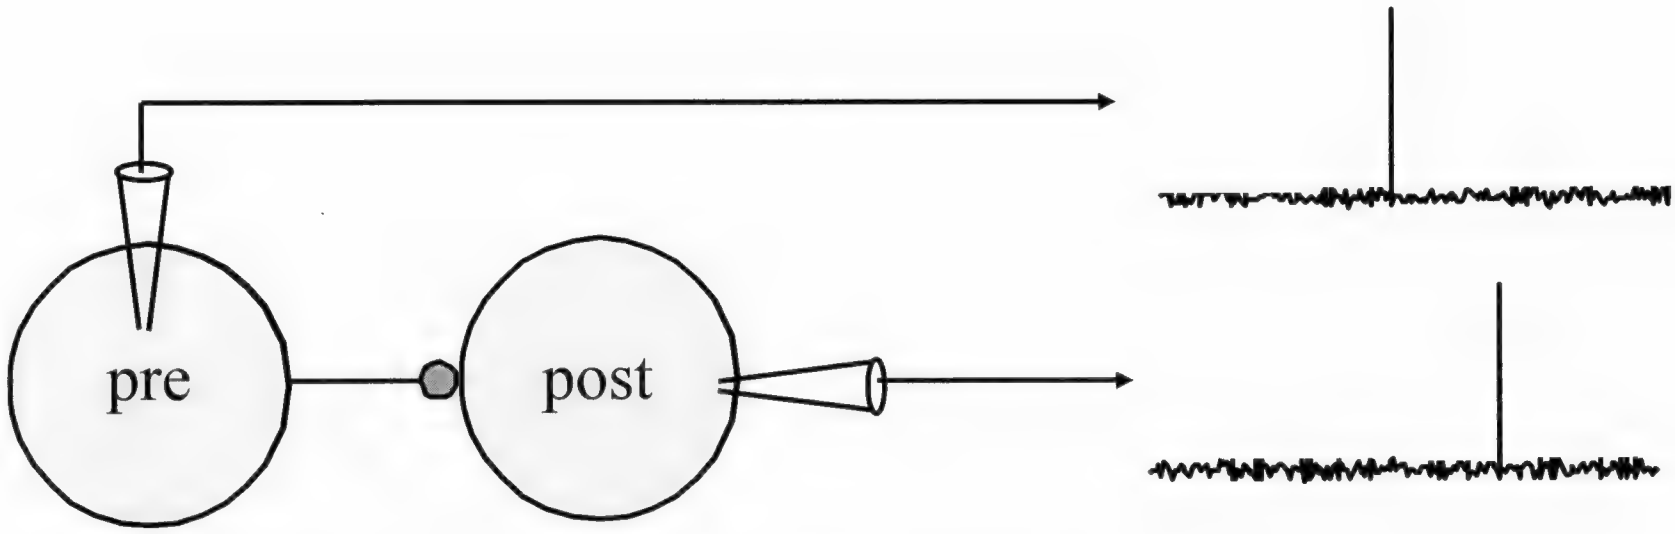
pre
post

$$\Delta w = F(\Delta t) = \begin{cases} -A_{-}\, e^{(\Delta t)/\tau_{-}} & \text{if } \Delta t \leq 0 \\ A_{+}\, e^{-(\Delta t)/\tau_{+}} & \text{if } \Delta t > 0 \end{cases}$$

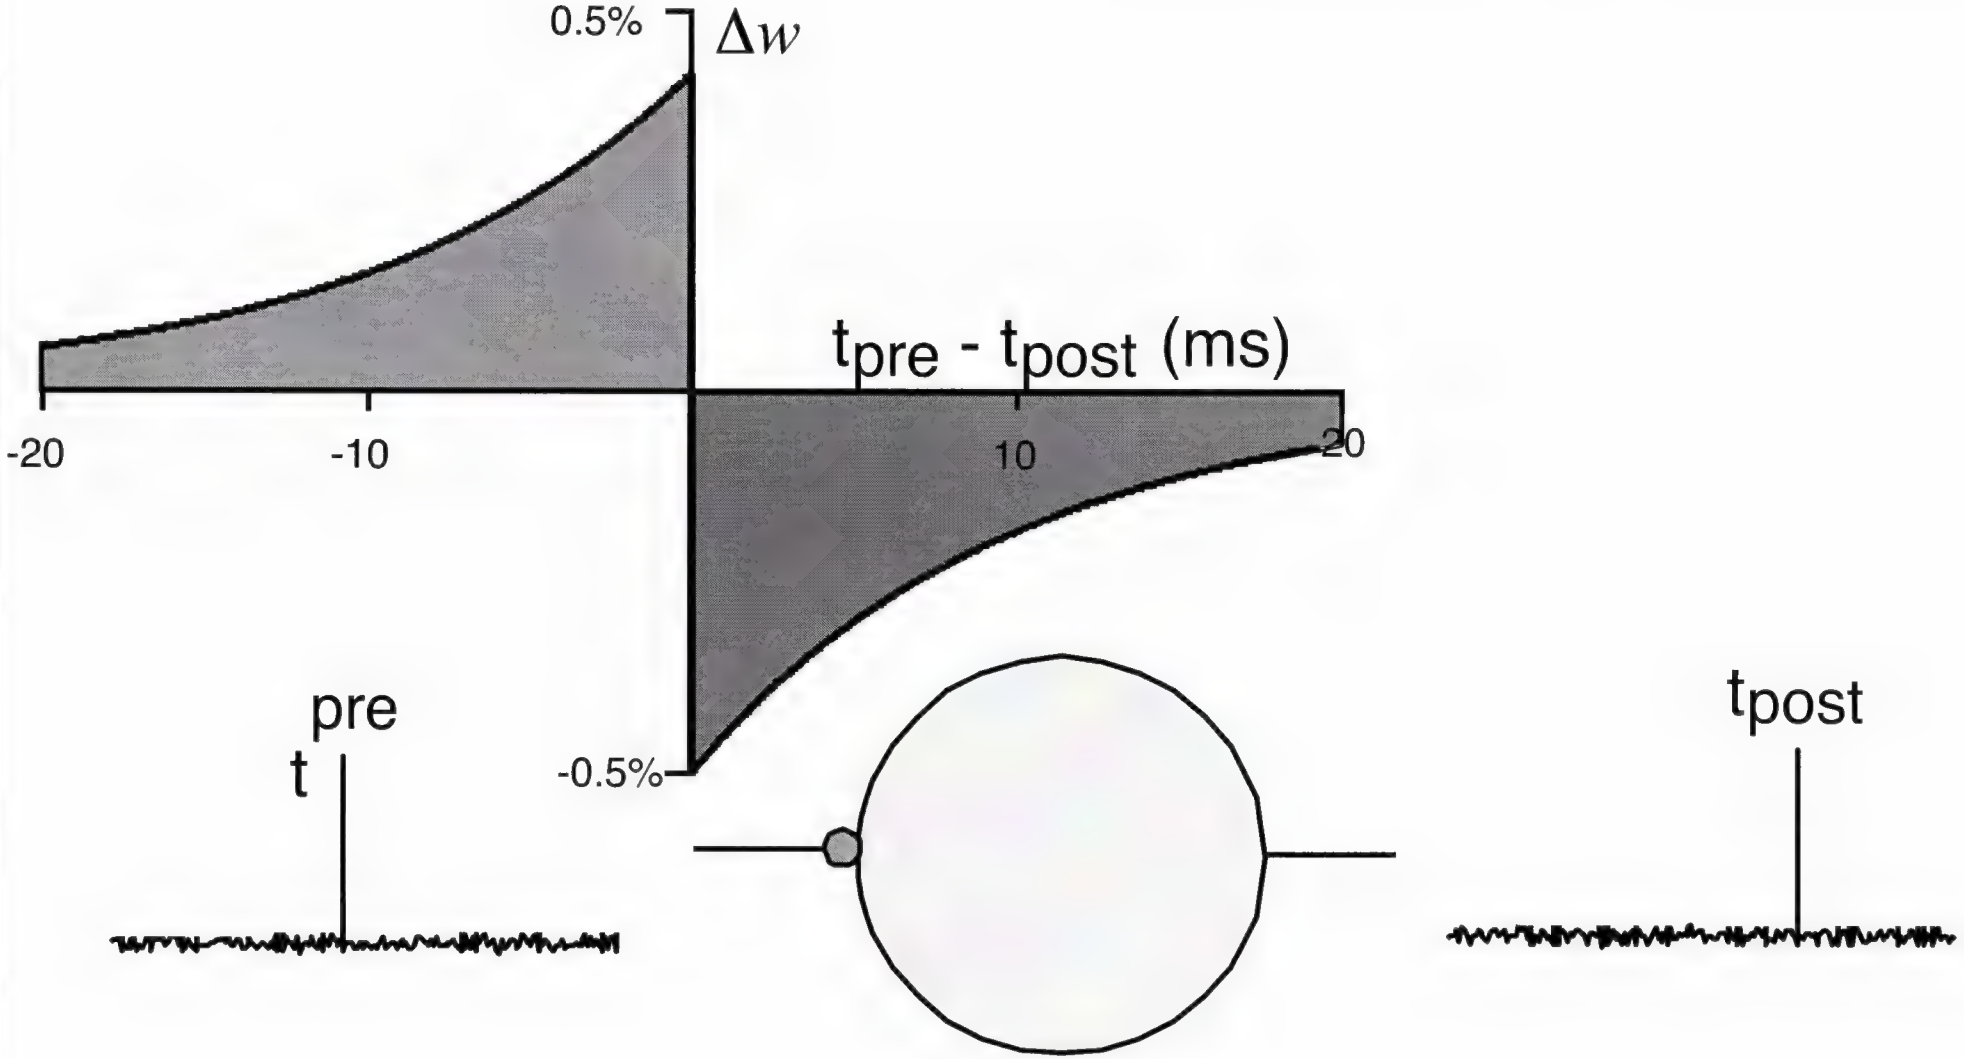

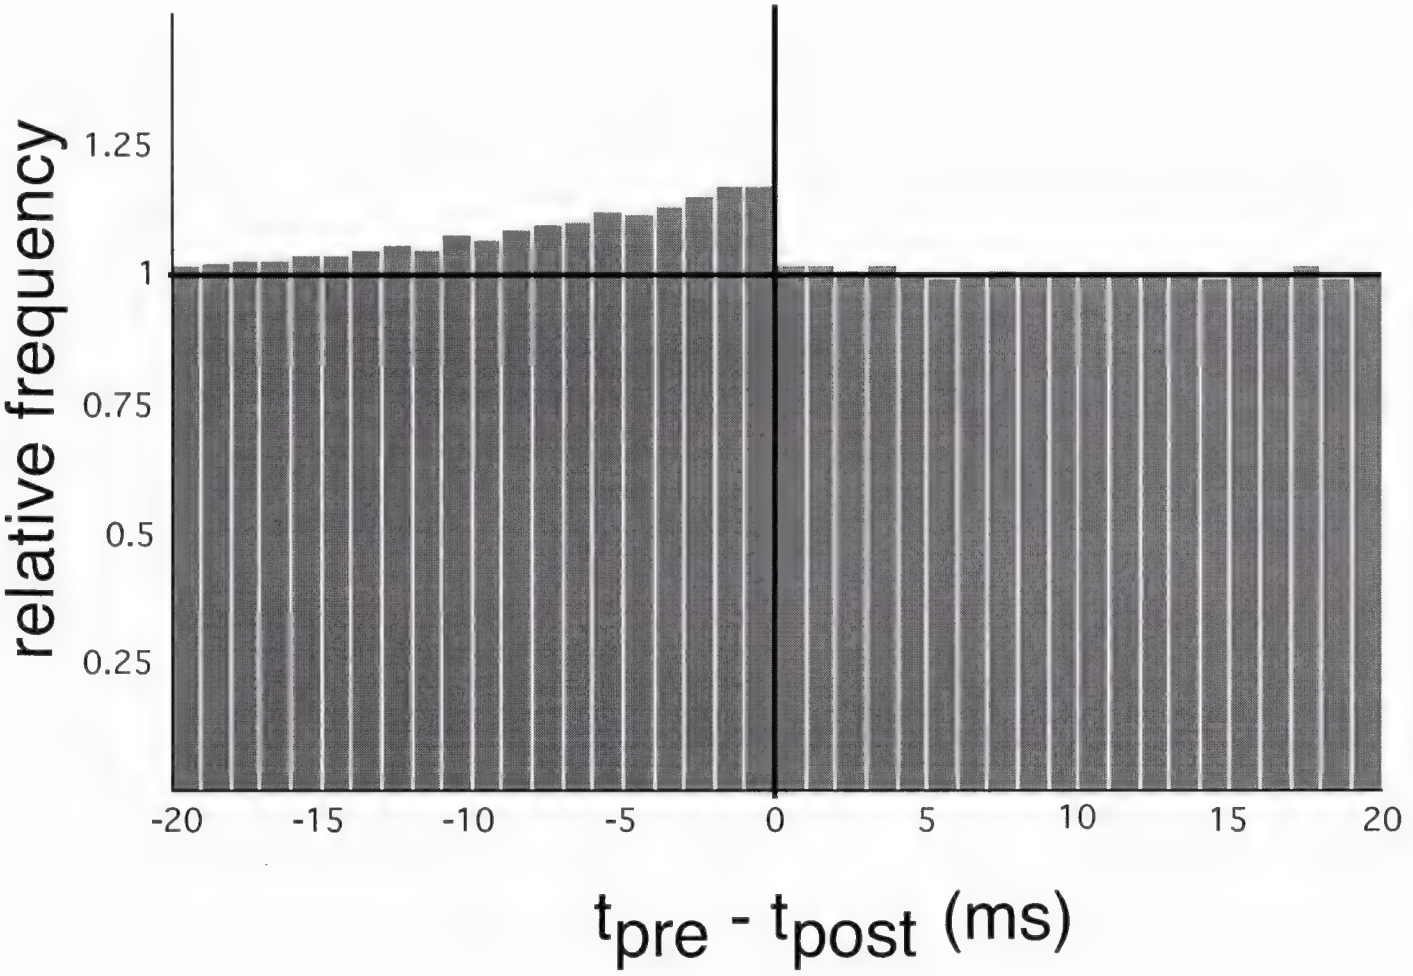
relative frequency
1.25
1
0.75
0.5
0.25
-20
-15
-10
-5
0
5
10
15
20
$t_{pre} - t_{post}$ (ms)

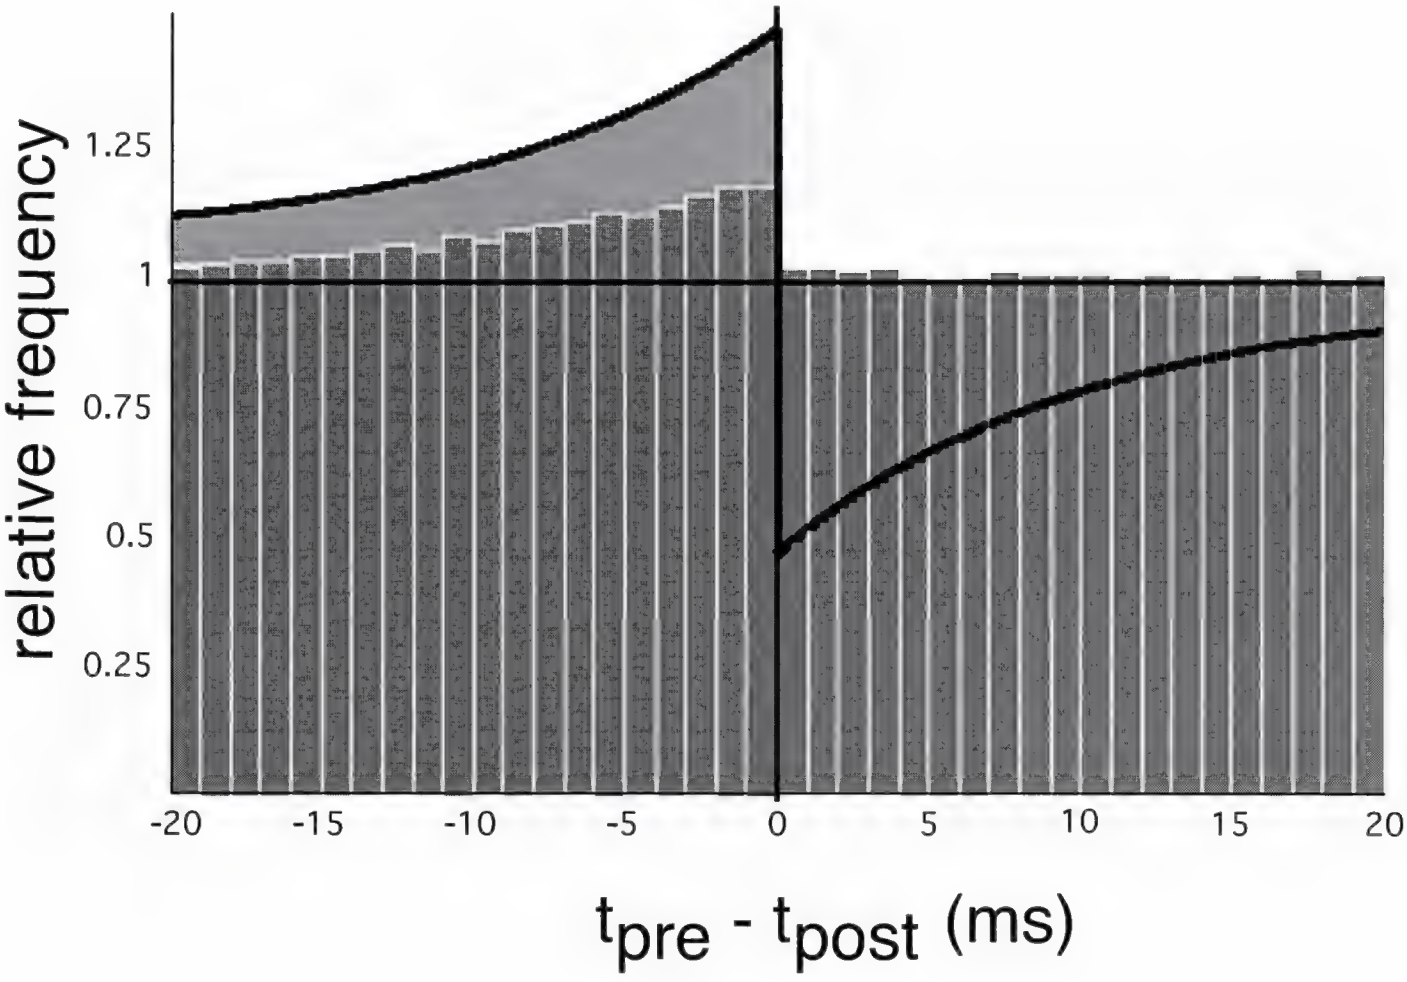
relative frequency
1.25
1
0.75
0.5
0.25
-20
-15
-10
-5
0
5
10
15
20
$t_{pre} - t_{post}$ (ms)

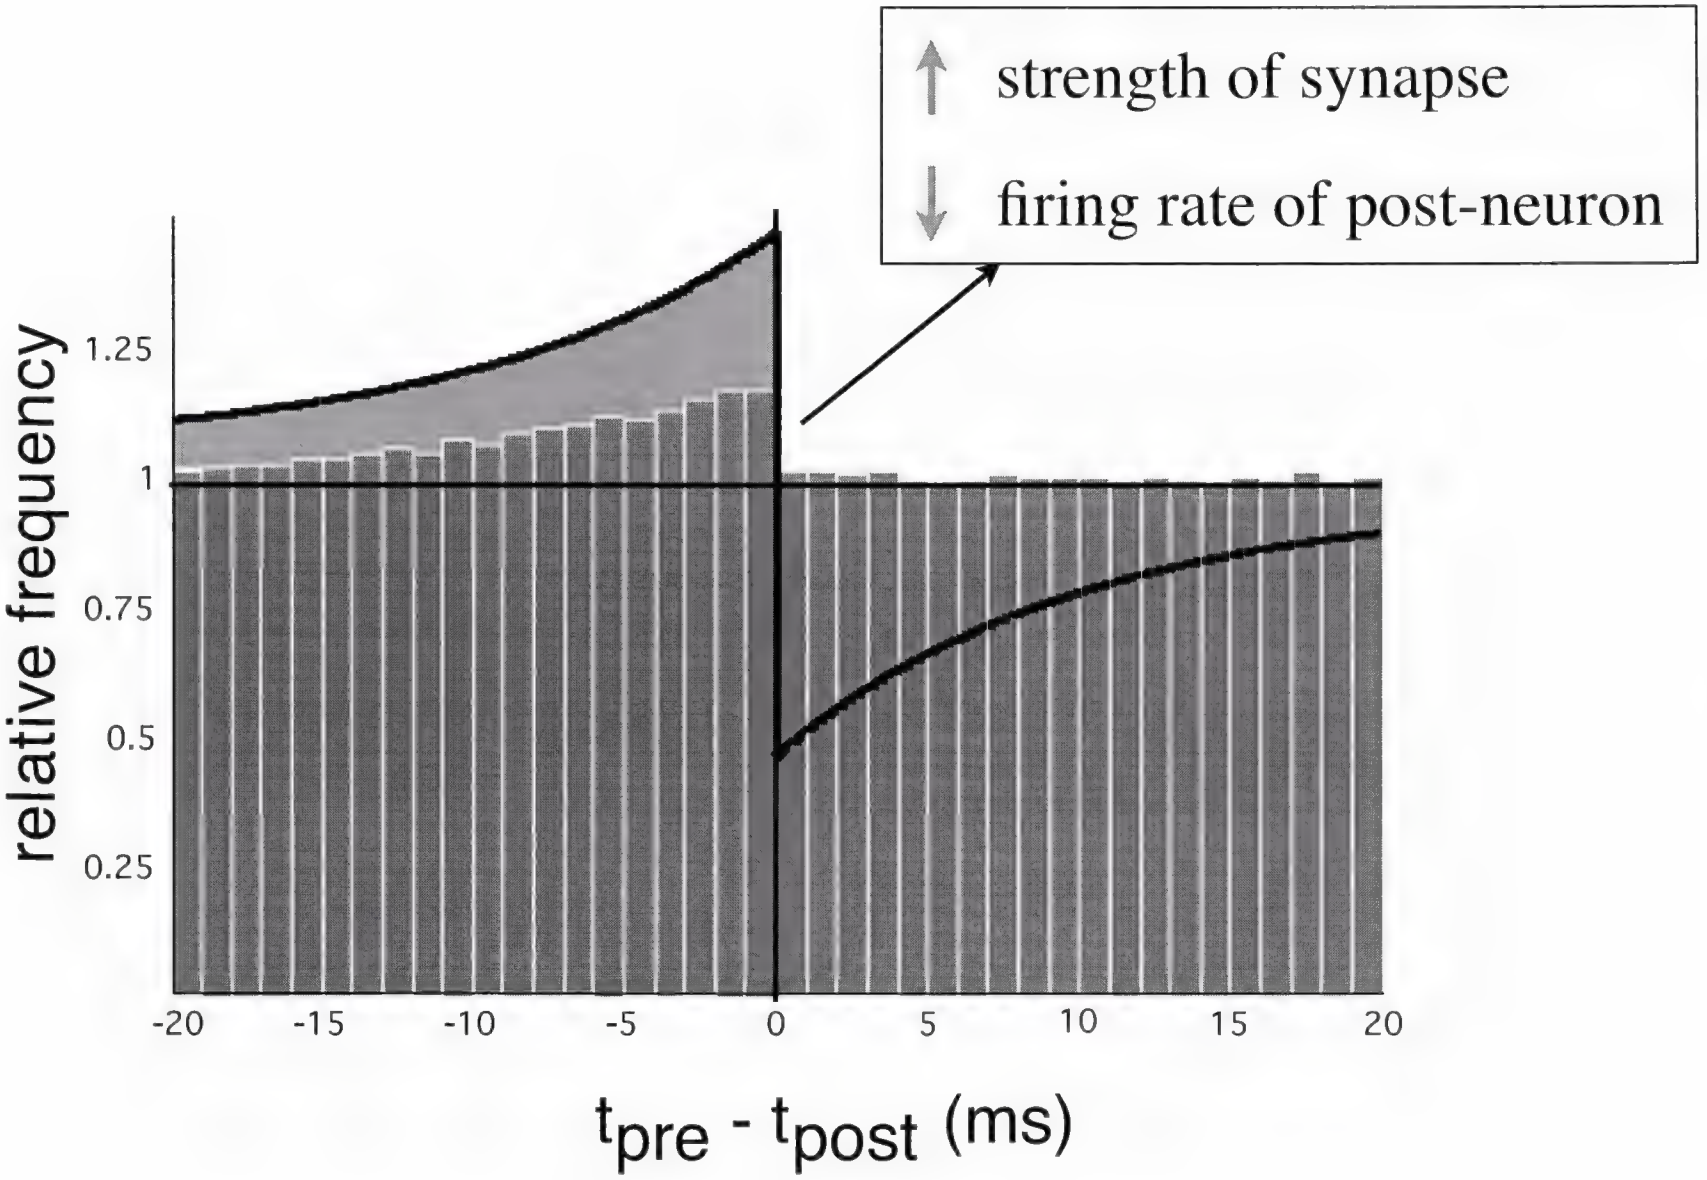
strength of synapse
firing rate of post-neuron
relative frequency
1.25
1
0.75
0.5
0.25
-20
-15
-10
-5
0
5
10
15
20
$t_{pre} - t_{post}$ (ms)

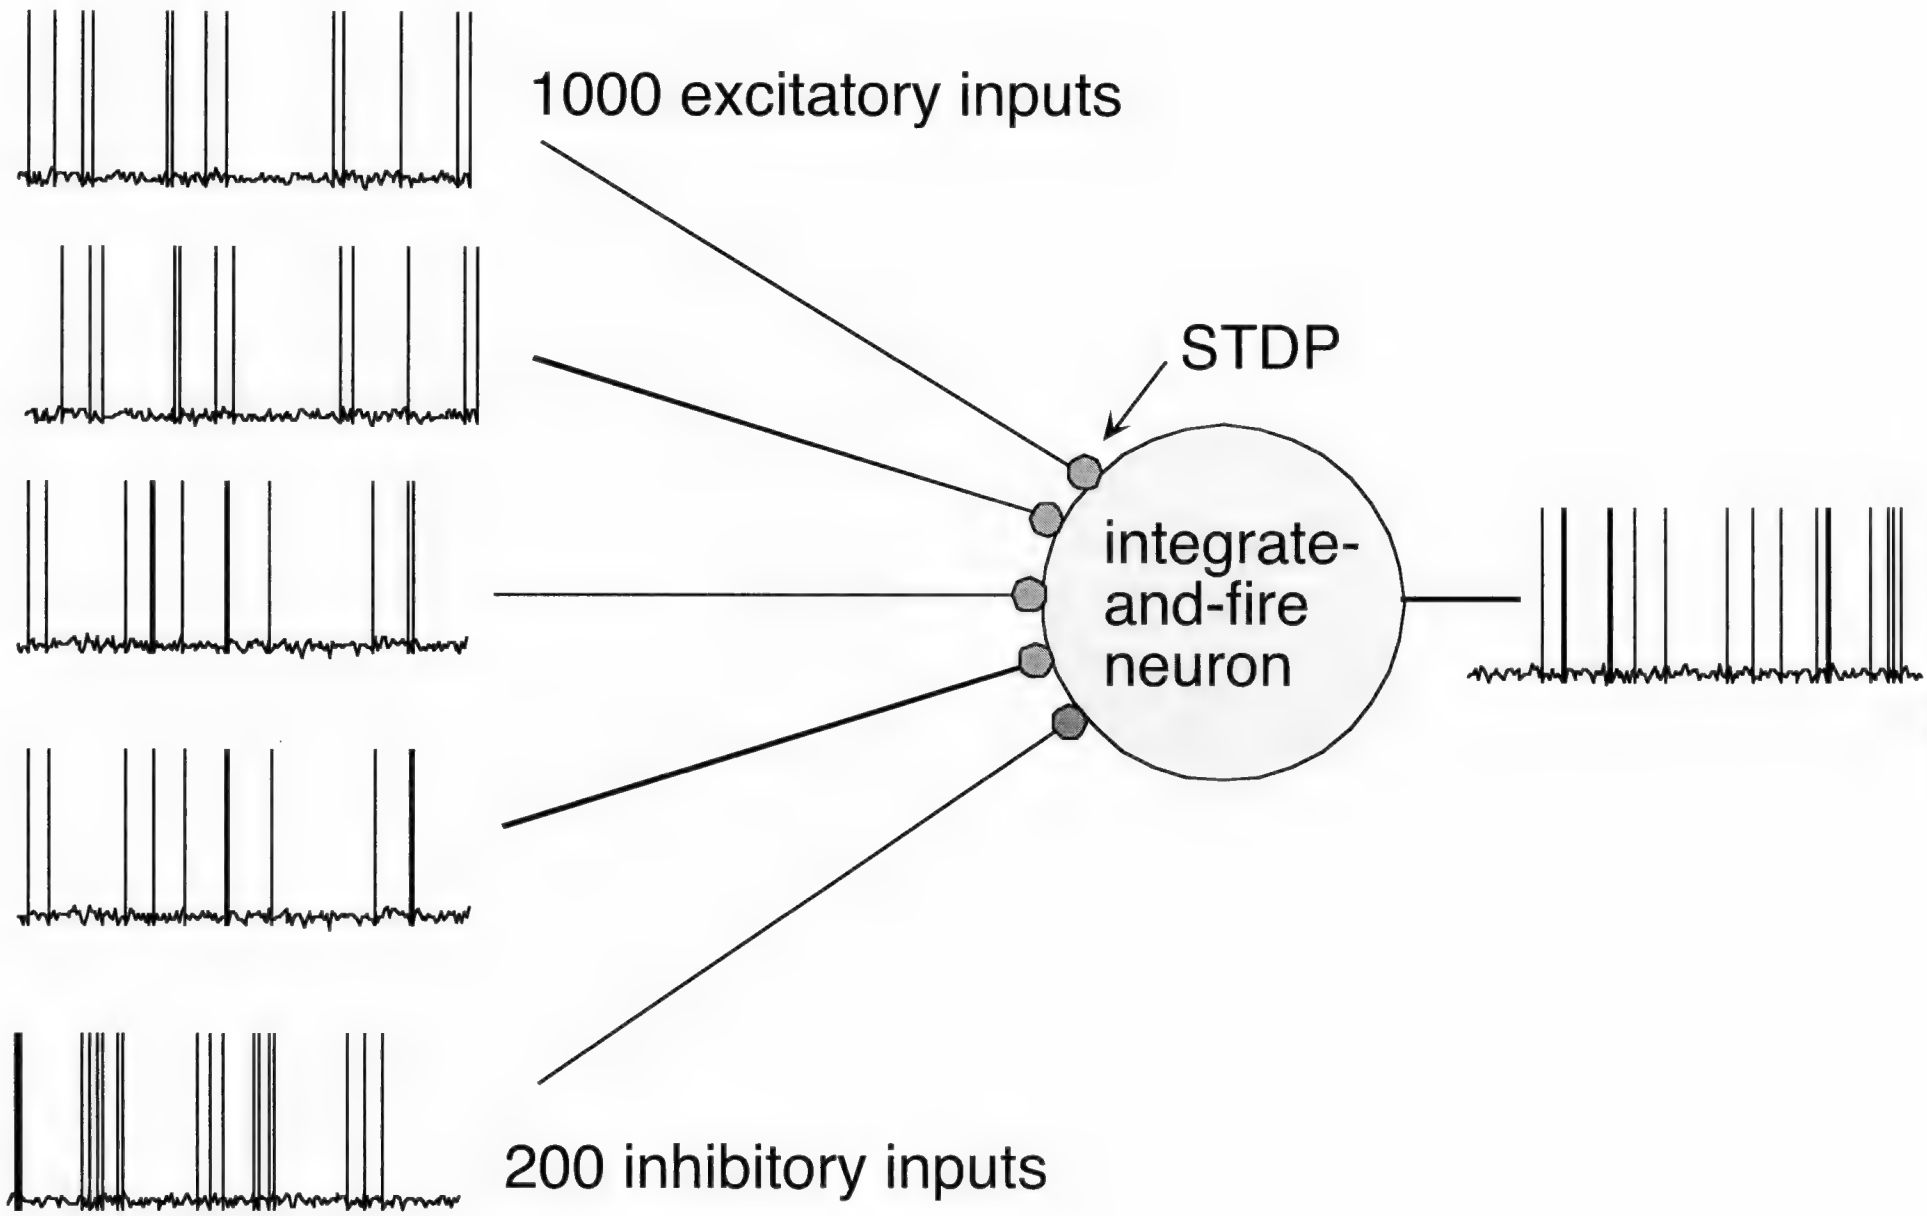
1000 excitatory inputs
STDP
integrate-
and-fire
neuron
200 inhibitory inputs

STDP Demo

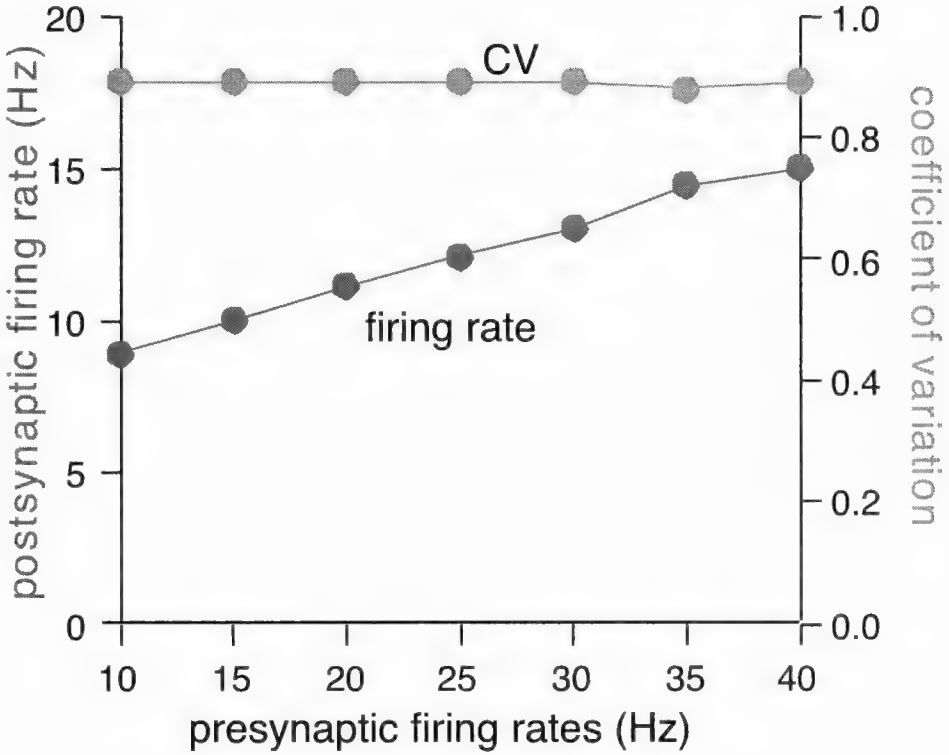
20
15
10
5
0
postsynaptic firing rate (Hz)
CV
firing rate
1.0
0.8
0.6
0.4
0.2
0.0
coefficient of variation
10
15
20
25
30
35
40
presynaptic firing rates (Hz)

STDP Correlation Demo

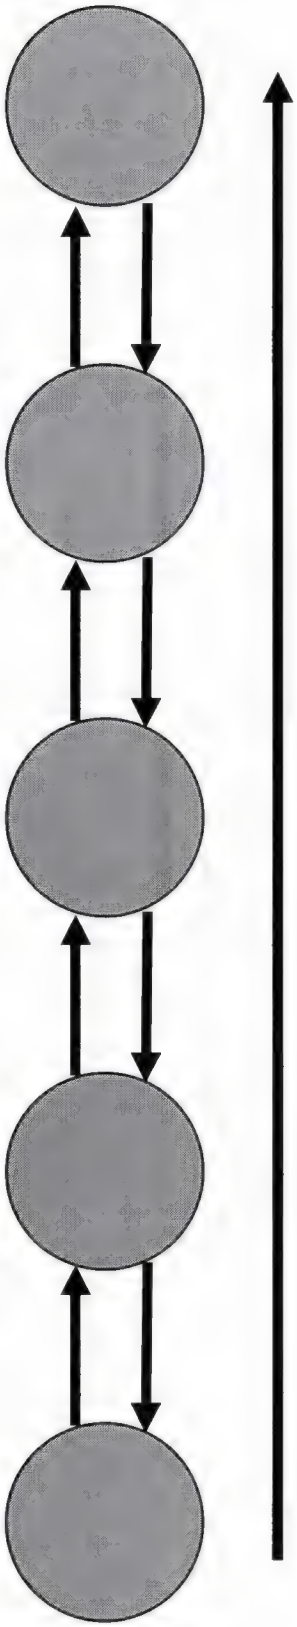

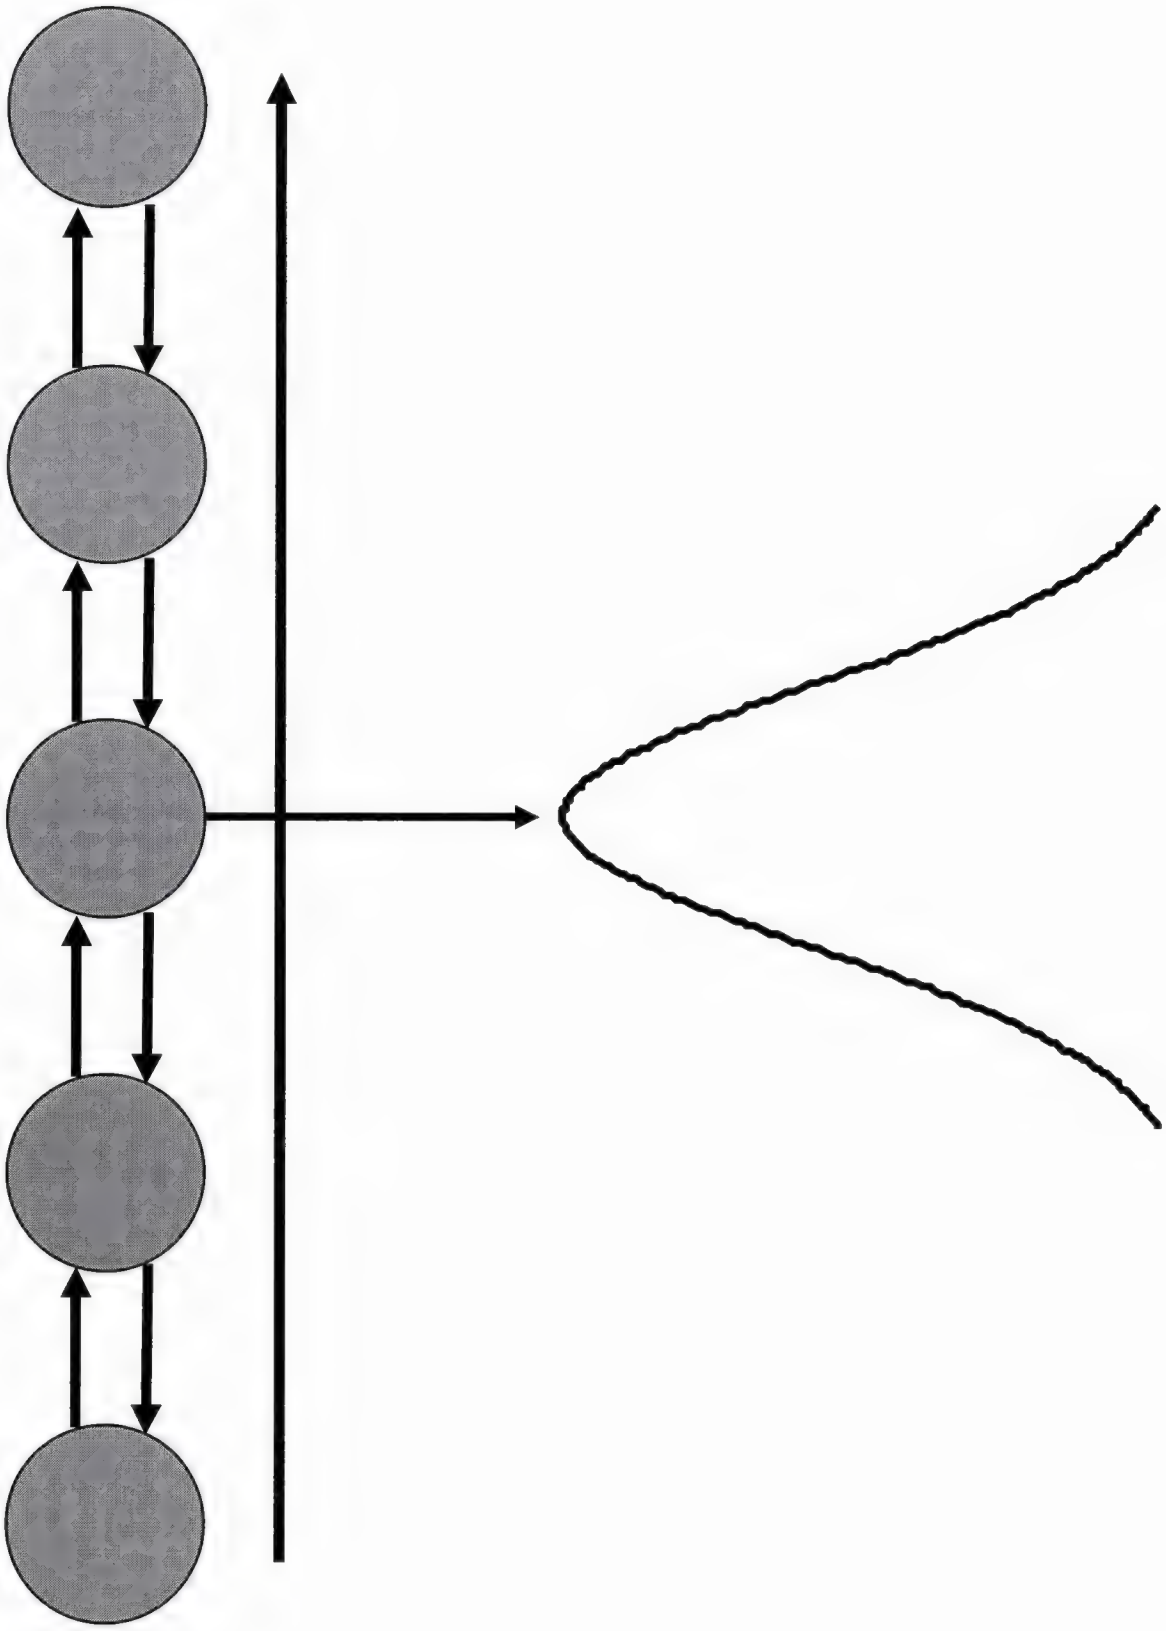

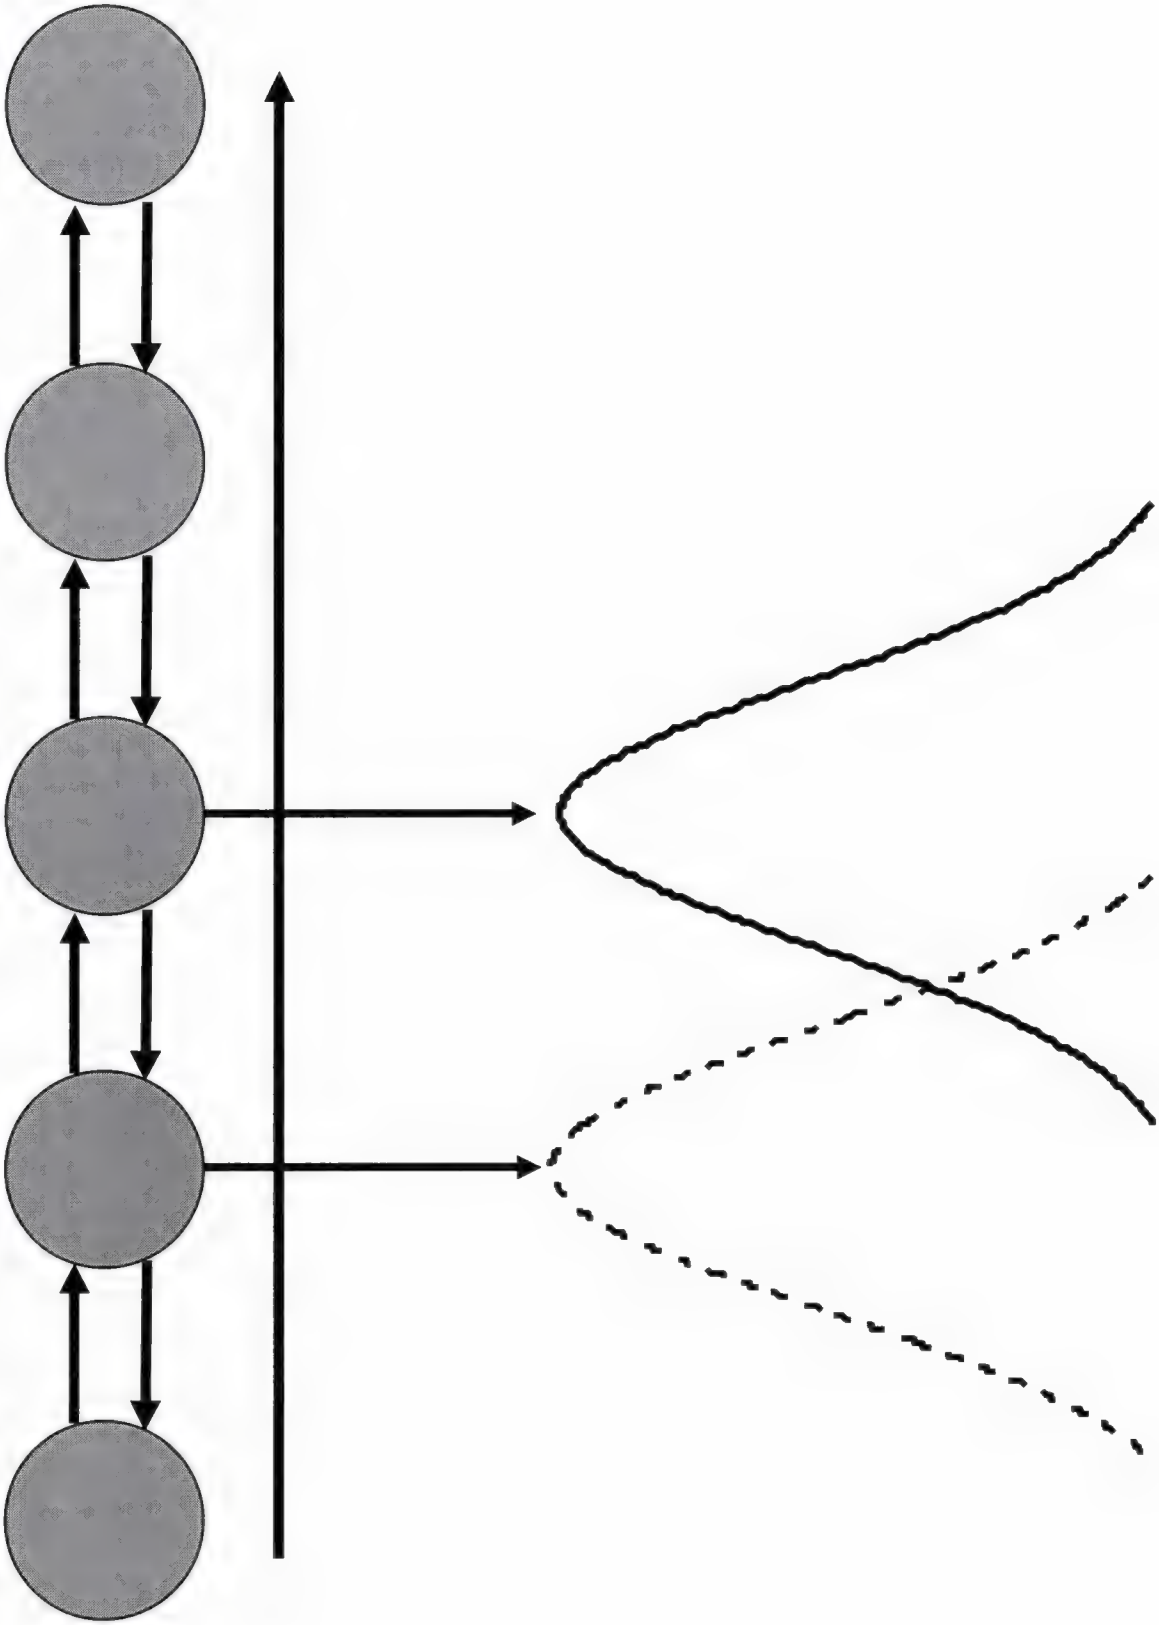

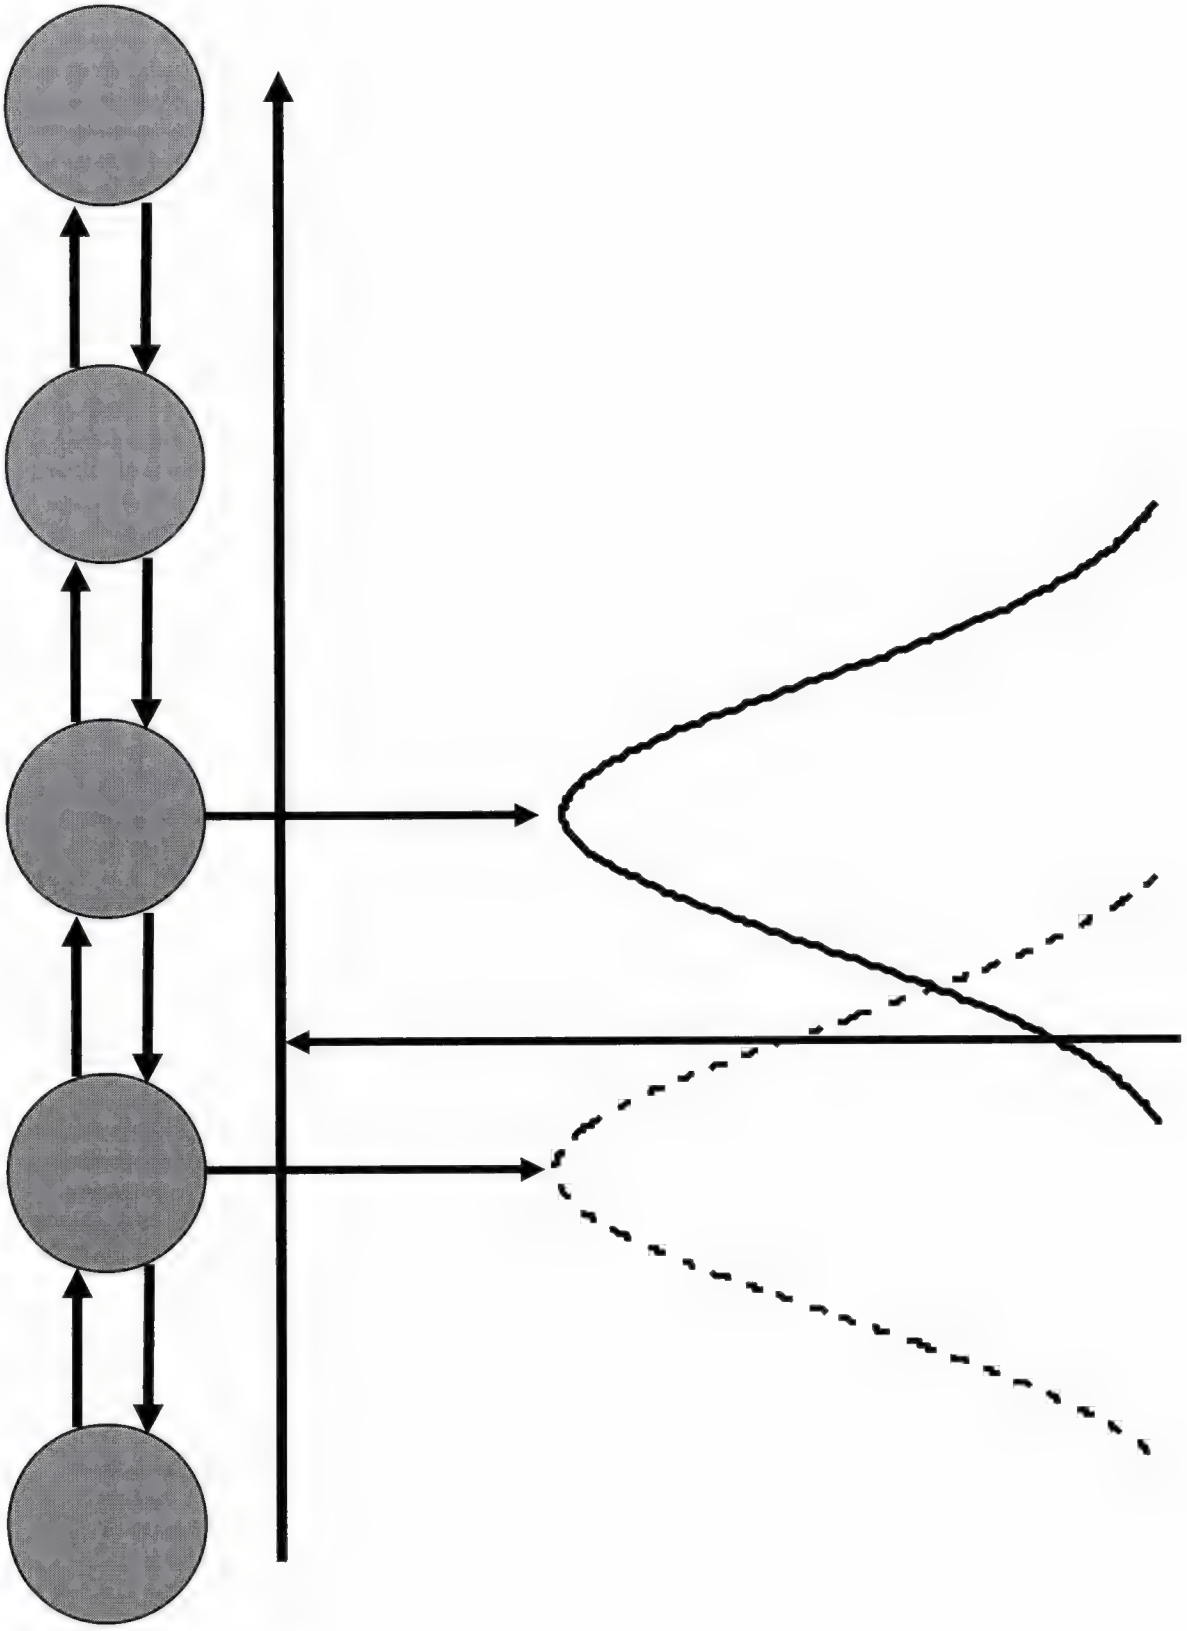

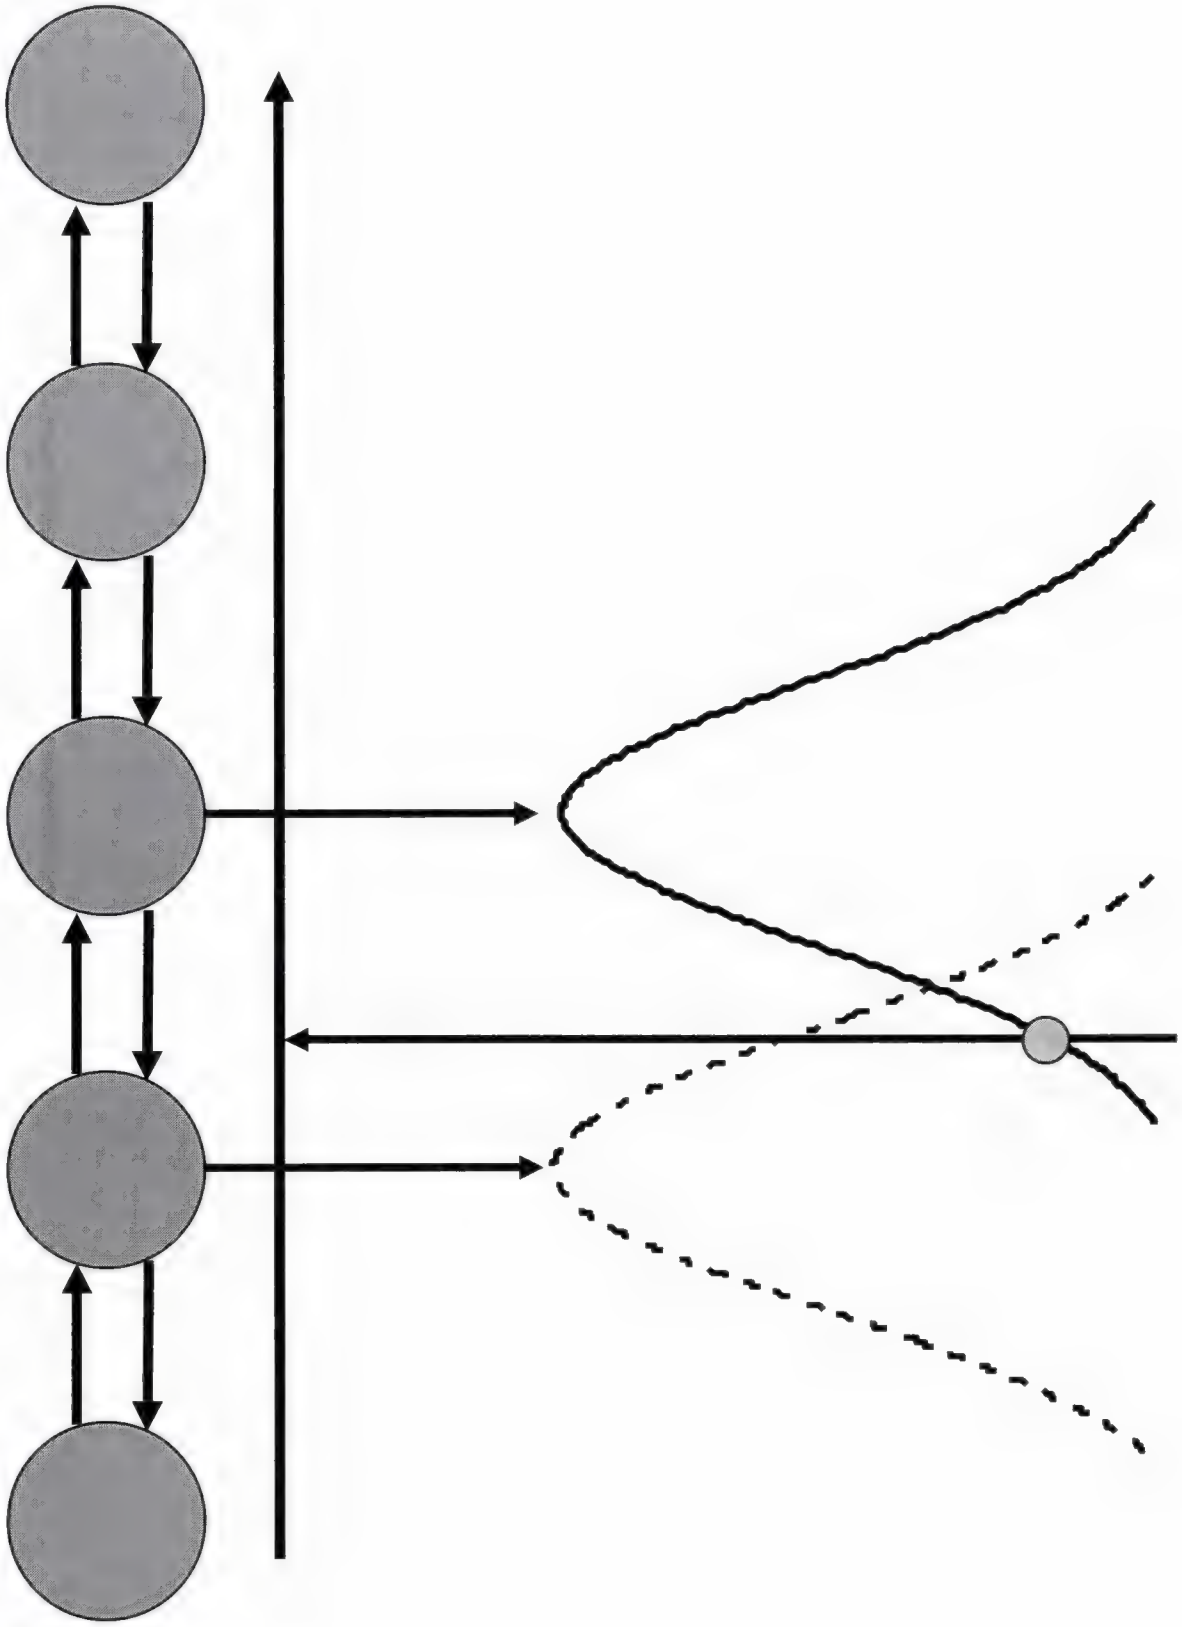

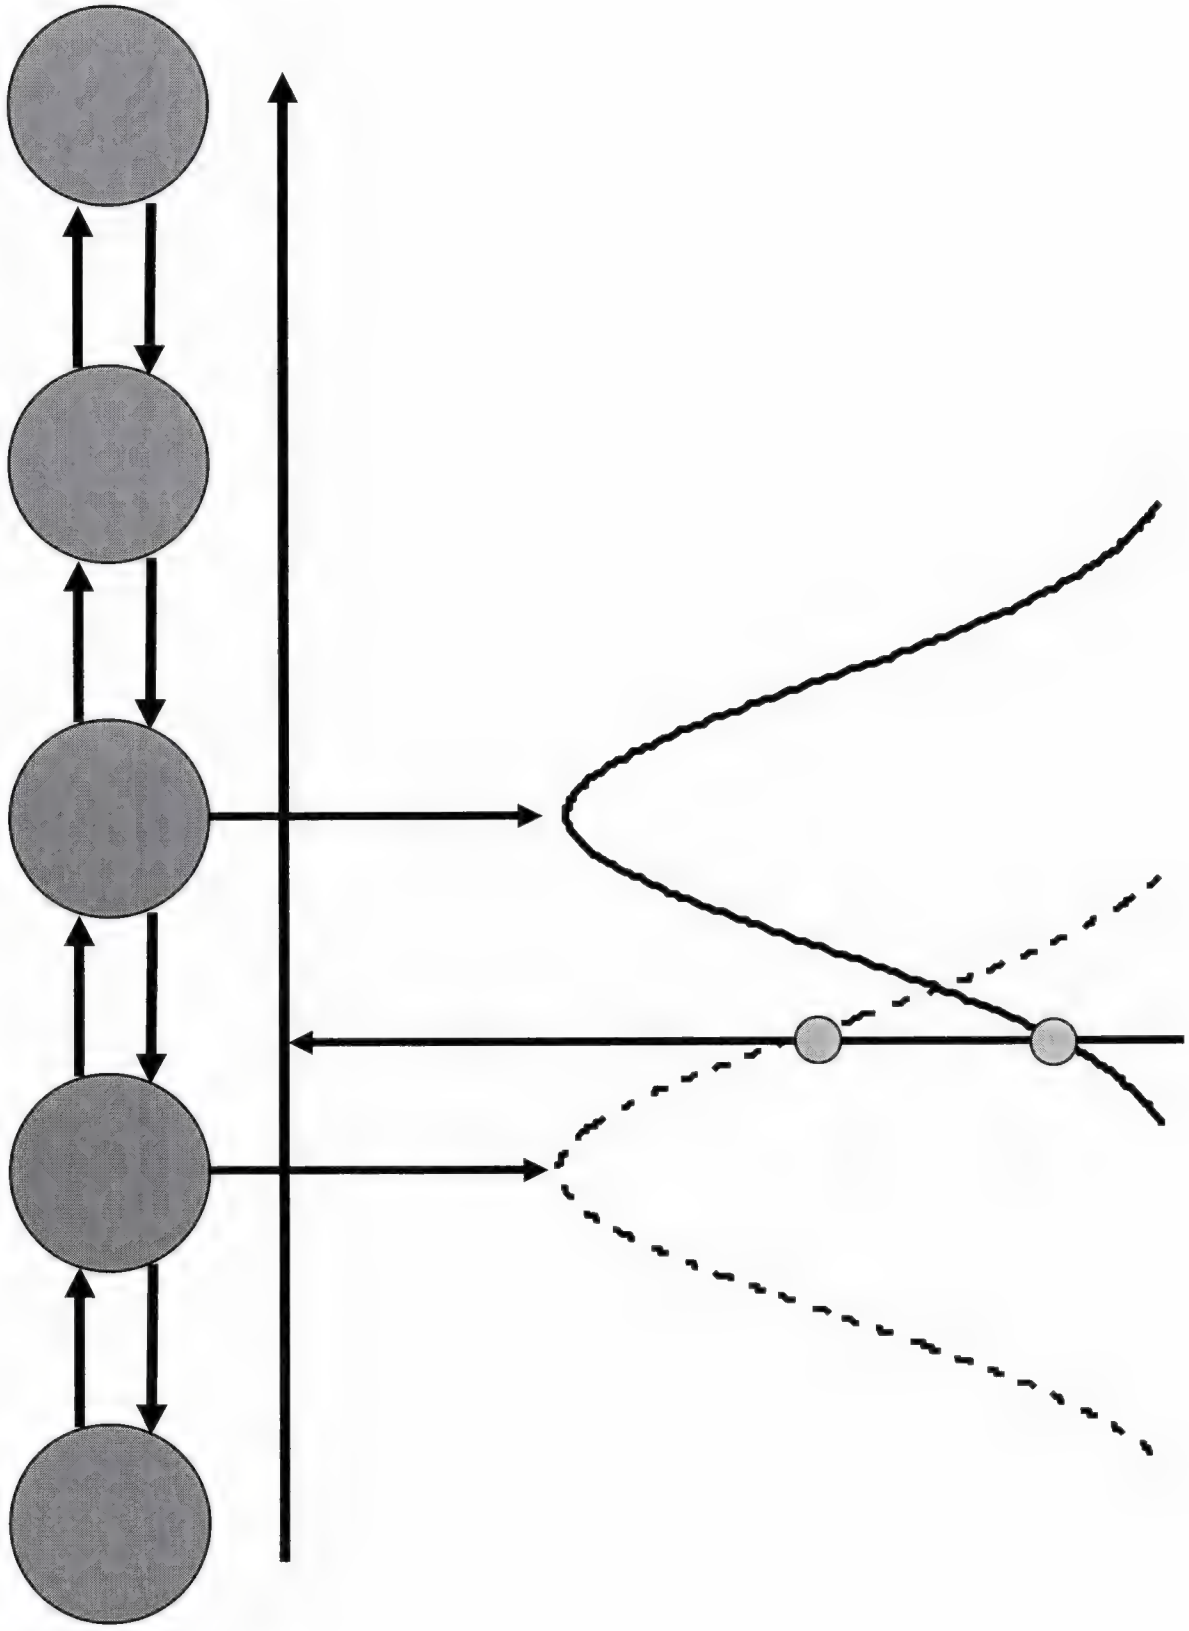

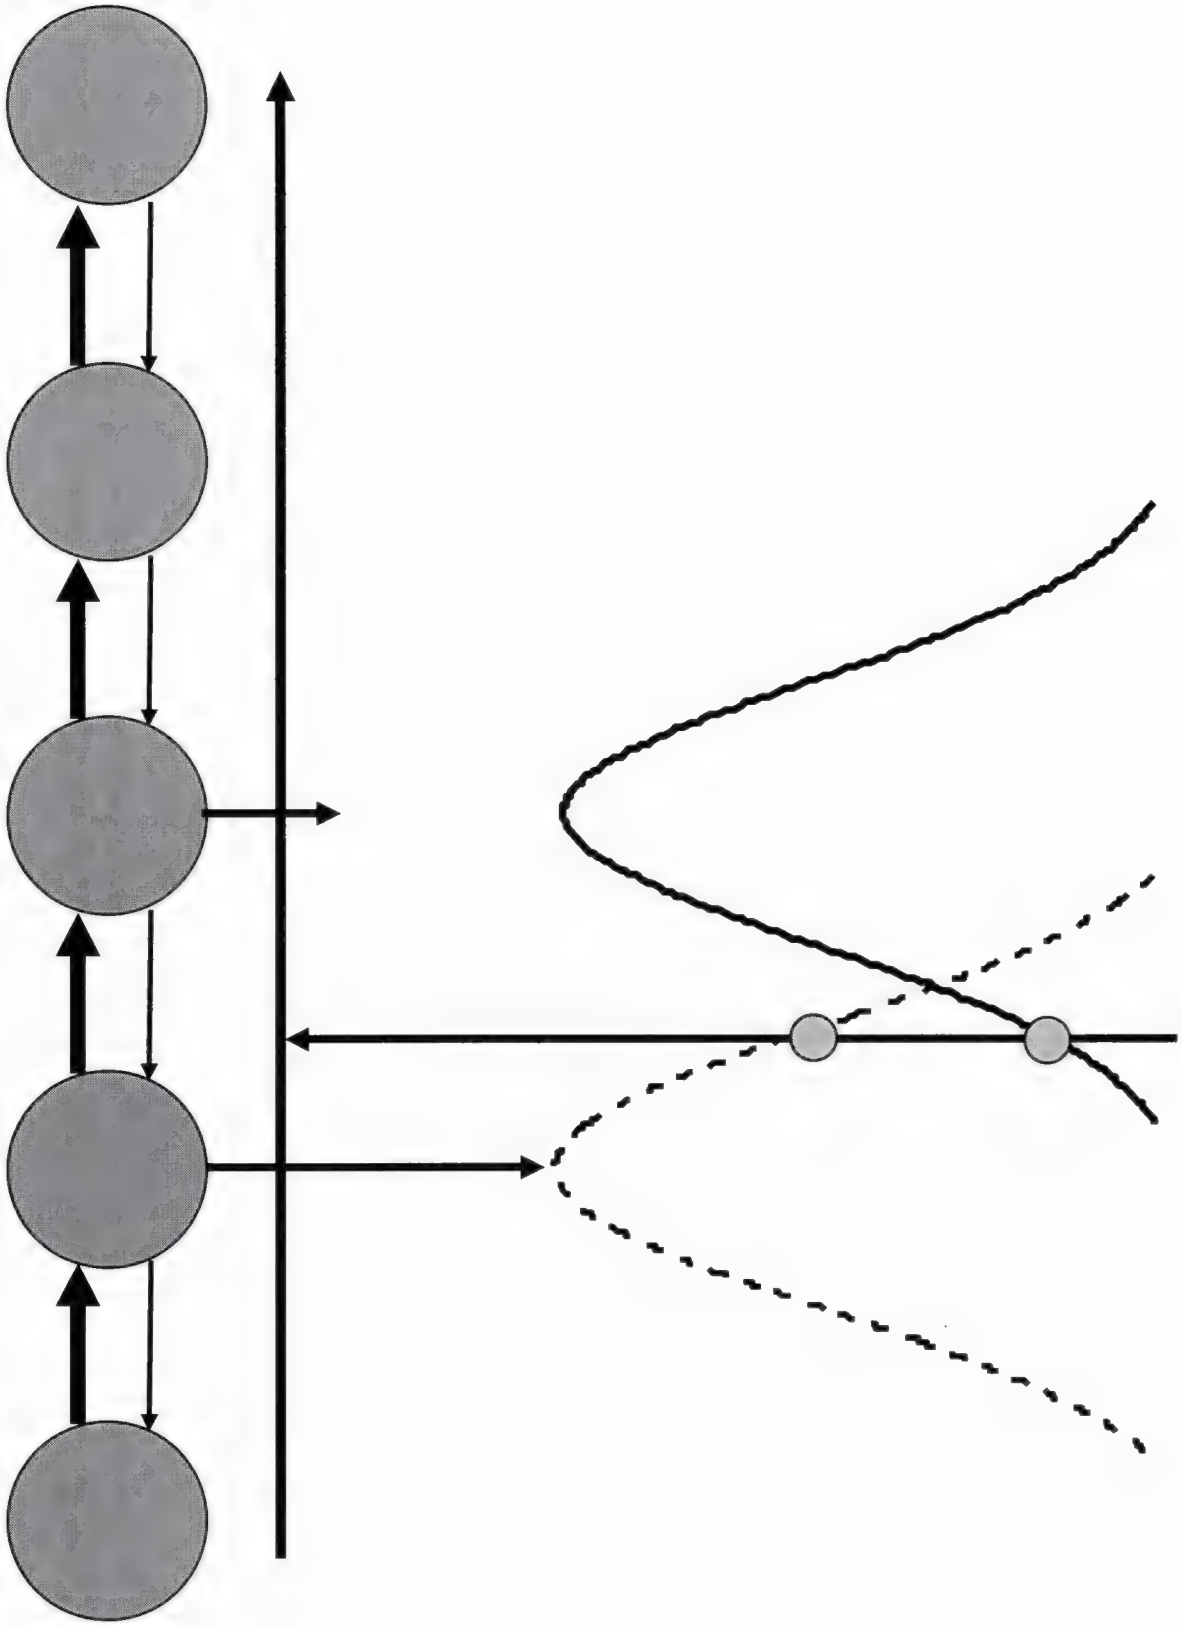

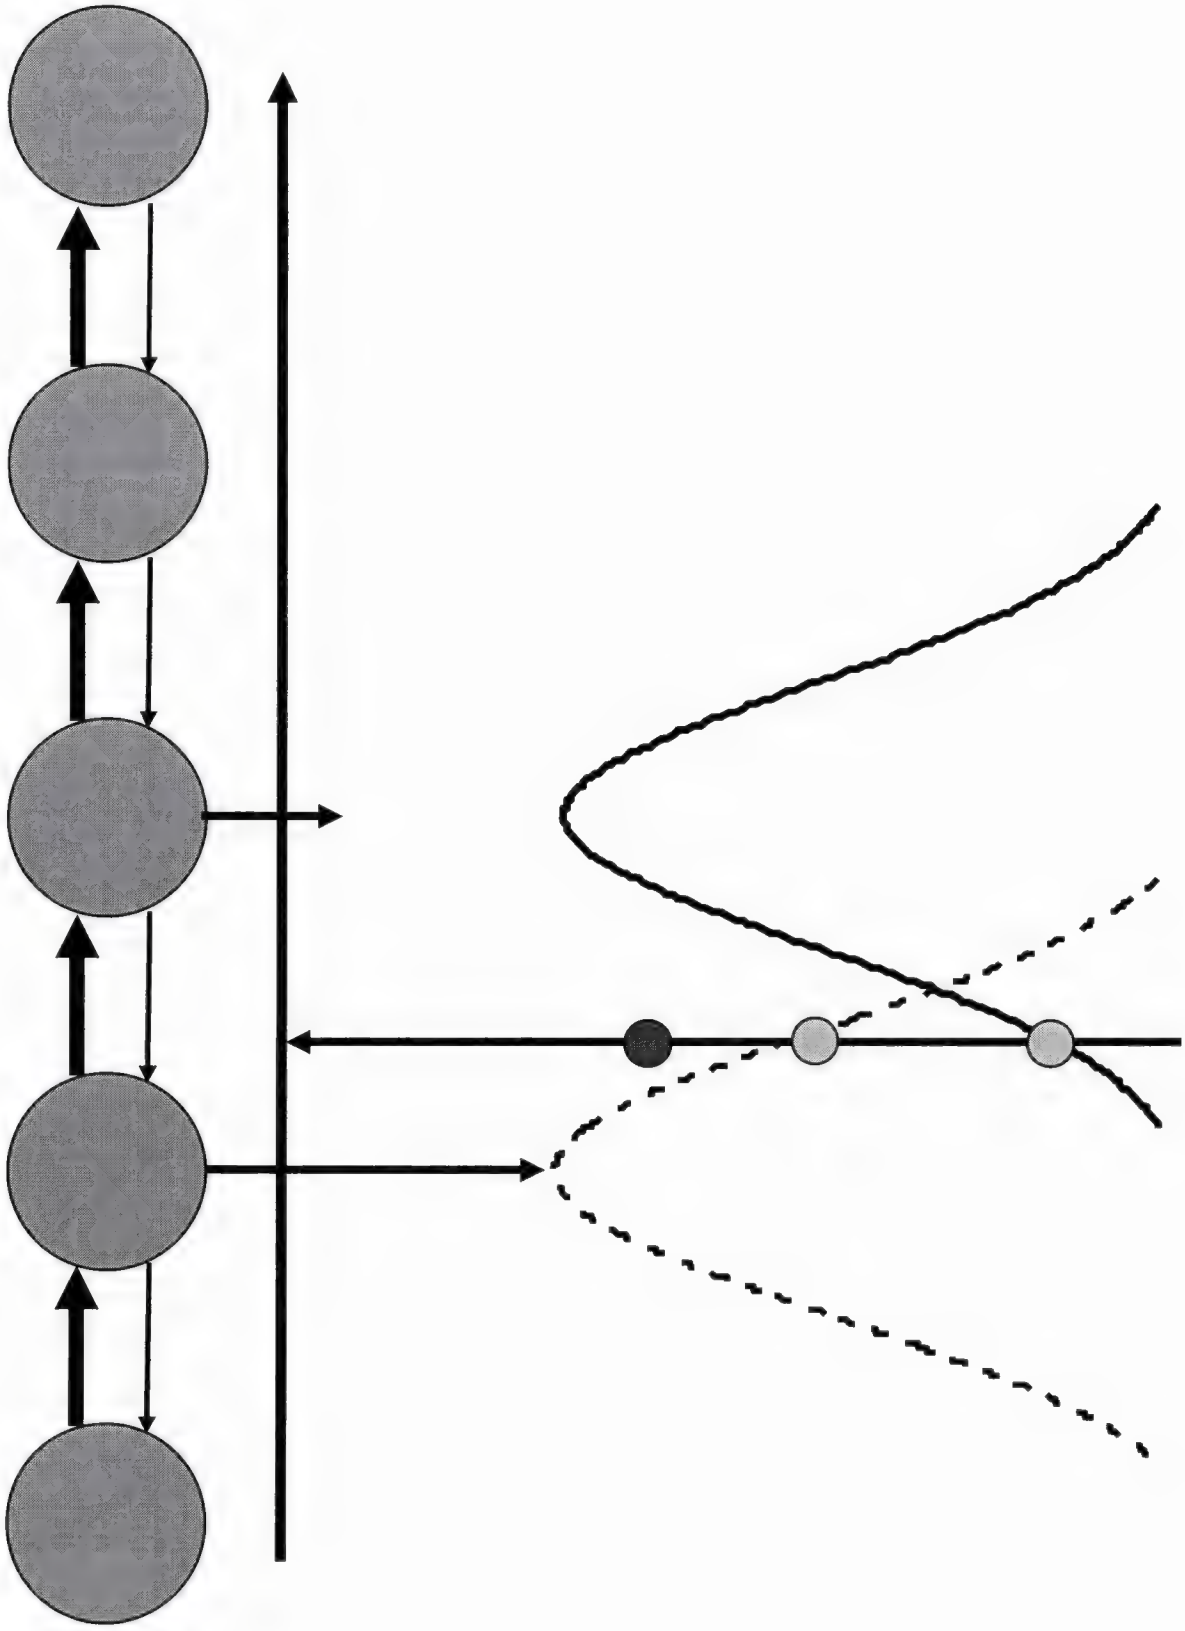

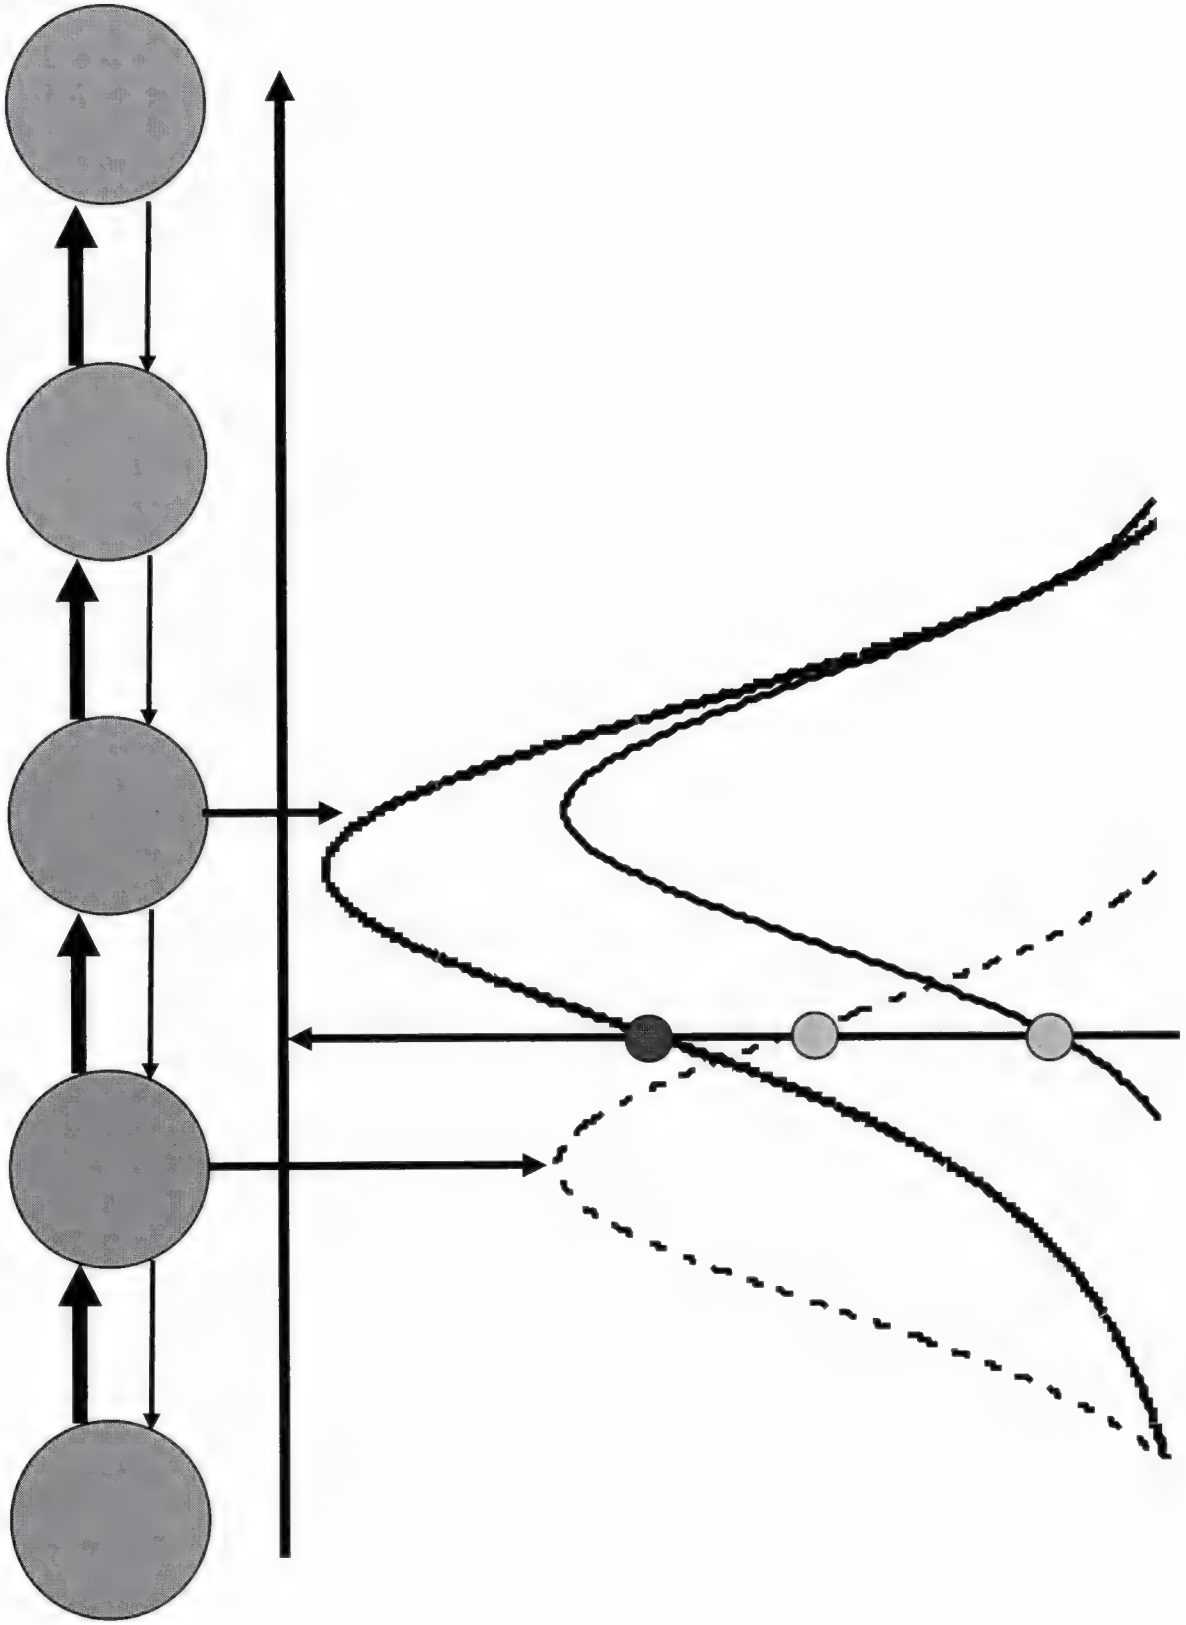

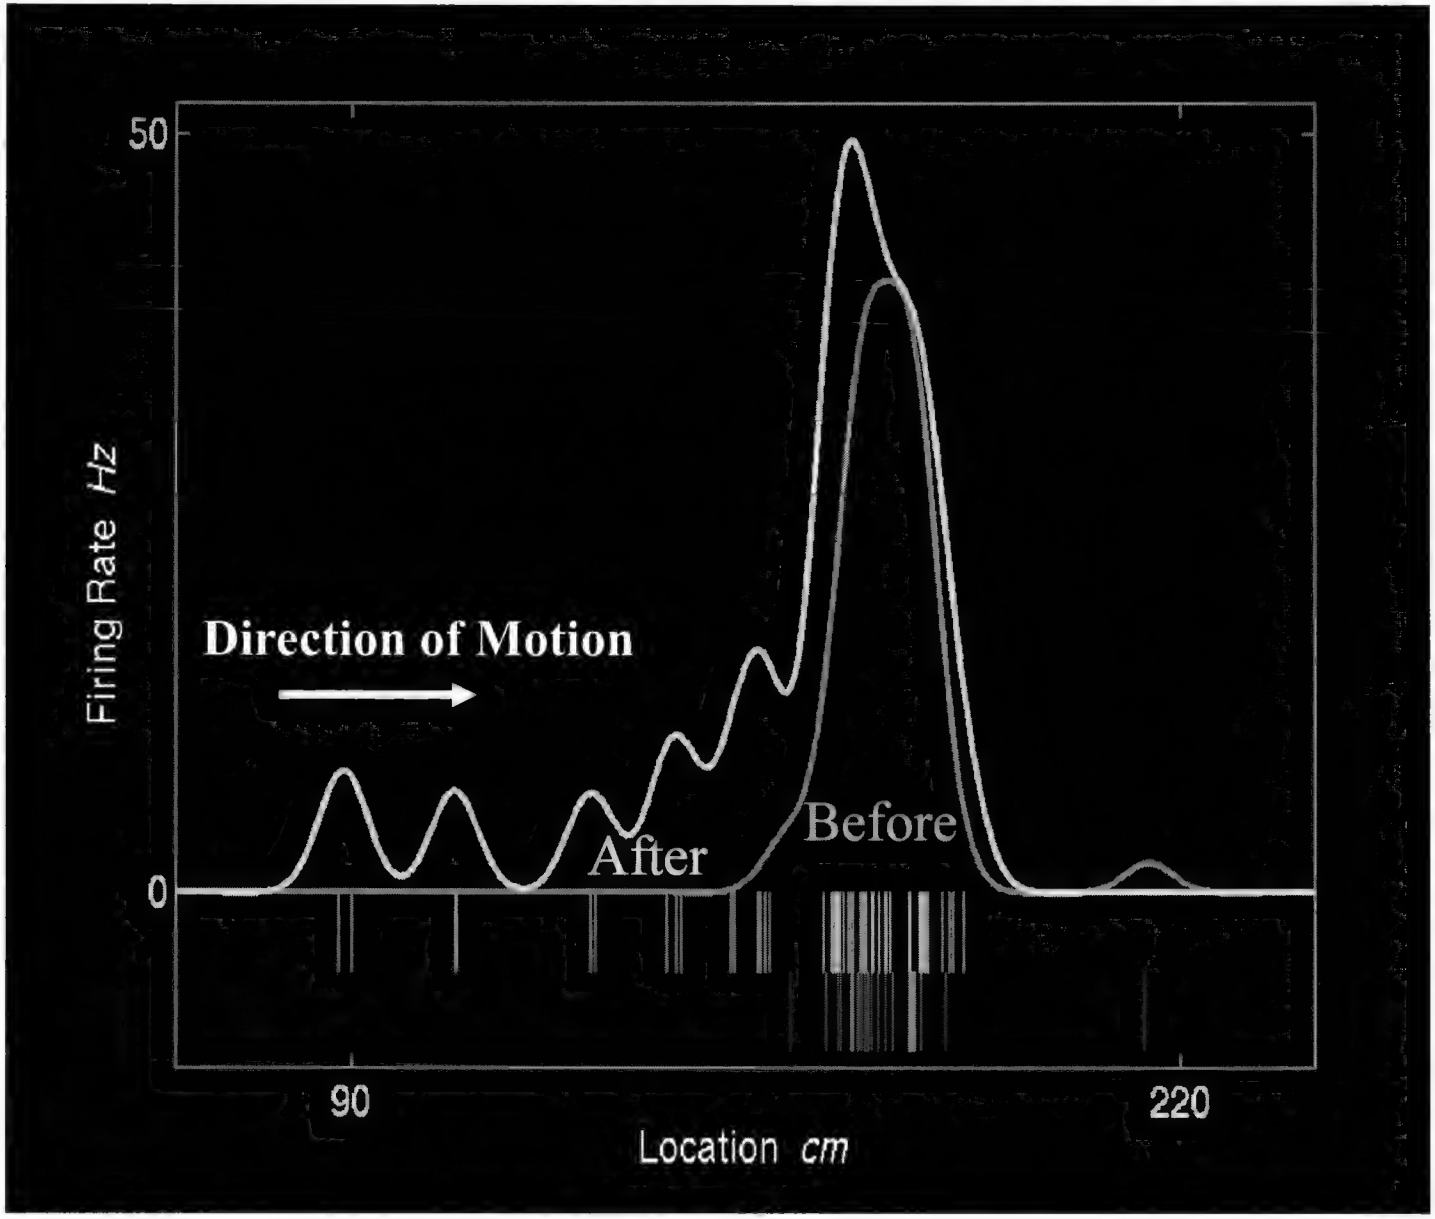

Mehta, Quirk & Wilson (2000) *Neuron* 25:707-715.

Mehta, Barnes, & McNaughton (1997) *PNAS* 94:8918-8921.

$$\Delta w = F(\Delta t) = \begin{cases} -f_-(w) A_- \, e^{(\Delta t)/\tau_-} & \text{if } \Delta t \leq 0 \\ \\ f_+(w) A_+ \, e^{-(\Delta t)/\tau_+} & \text{if } \Delta t > 0 \end{cases}$$

$$f_-(w) = w^{\mu} \quad \text{and} \quad f_+(w) = (1 - w)^{\mu}$$

Gutig, Aharanov, Rotter, Sompolinsky, 2003

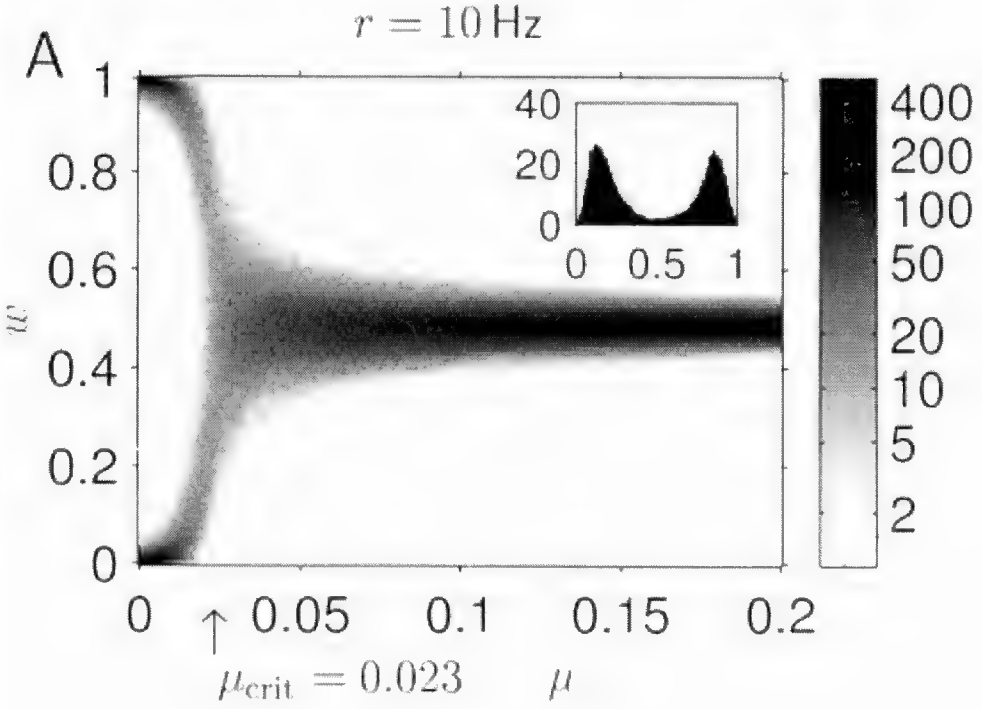
A
$r = 10\,\mathrm{Hz}$
1
0.8
0.6
0.4
0.2
0
w
40
20
0
0 0.5 1
400
200
100
50
20
10
5
2
0 0.05 0.1 0.15 0.2
$\mu_{\mathrm{crit}} = 0.023$
μ

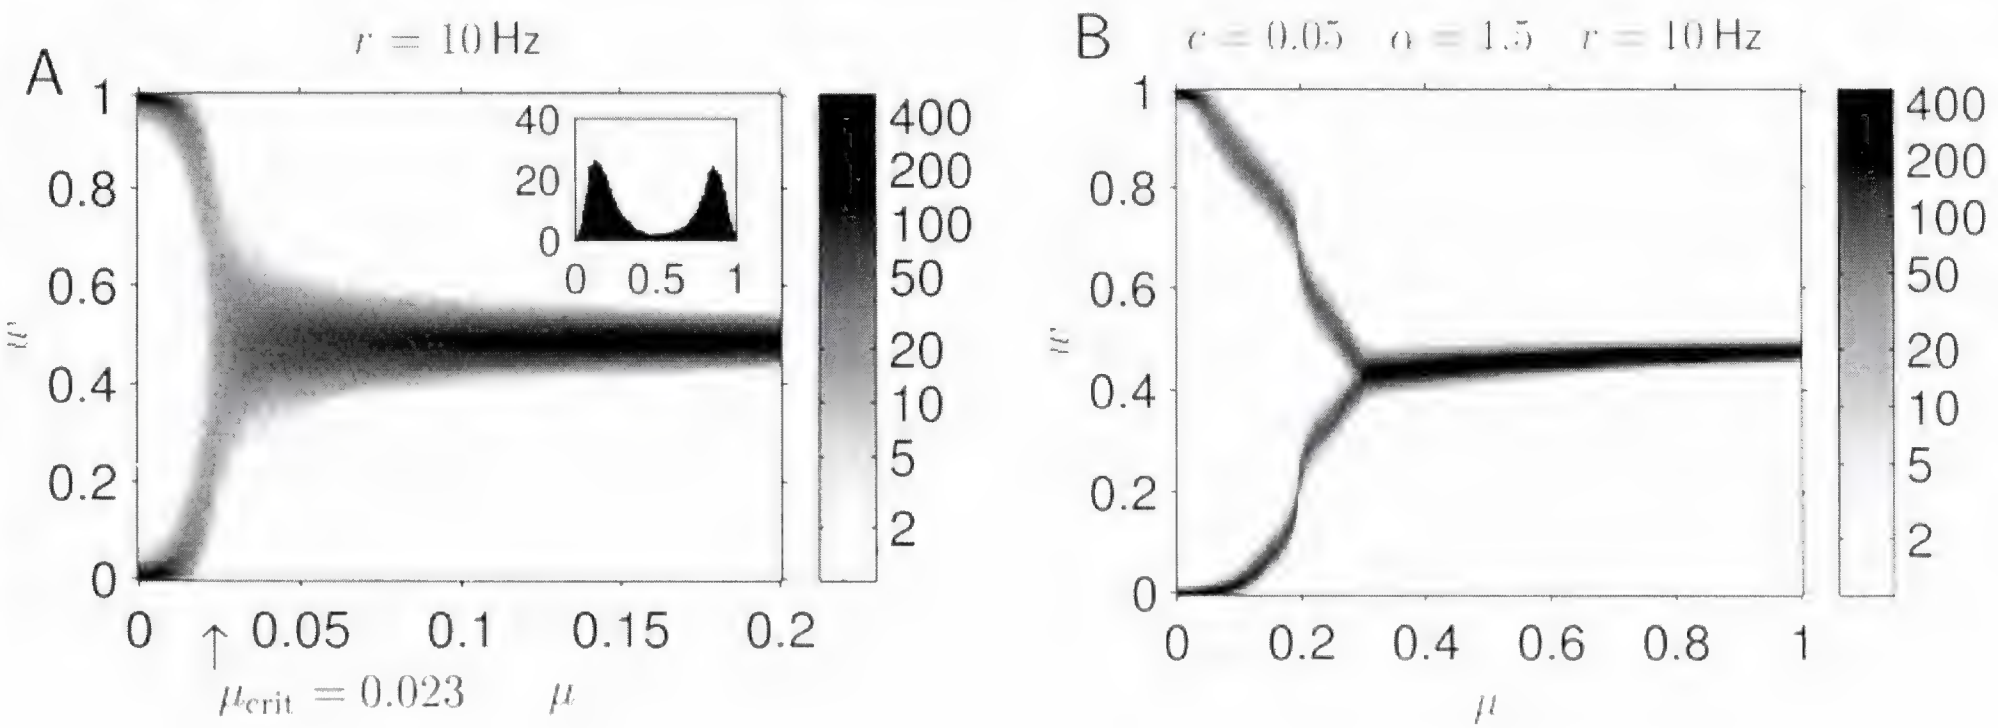
A
$r = 10\,\mathrm{Hz}$
1
0.8
0.6
0.4
0.2
0
w
40
20
0
0 0.5 1
400
200
100
50
20
10
5
2
0 0.05 0.1 0.15 0.2
$\mu_{\mathrm{crit}} = 0.023$
μ
B
$c = 0.05$ $\alpha = 1.5$ $r = 10\,\mathrm{Hz}$
1
0.8
0.6
0.4
0.2
0
w
400
200
100
50
20
10
5
2
0 0.2 0.4 0.6 0.8 1
μ

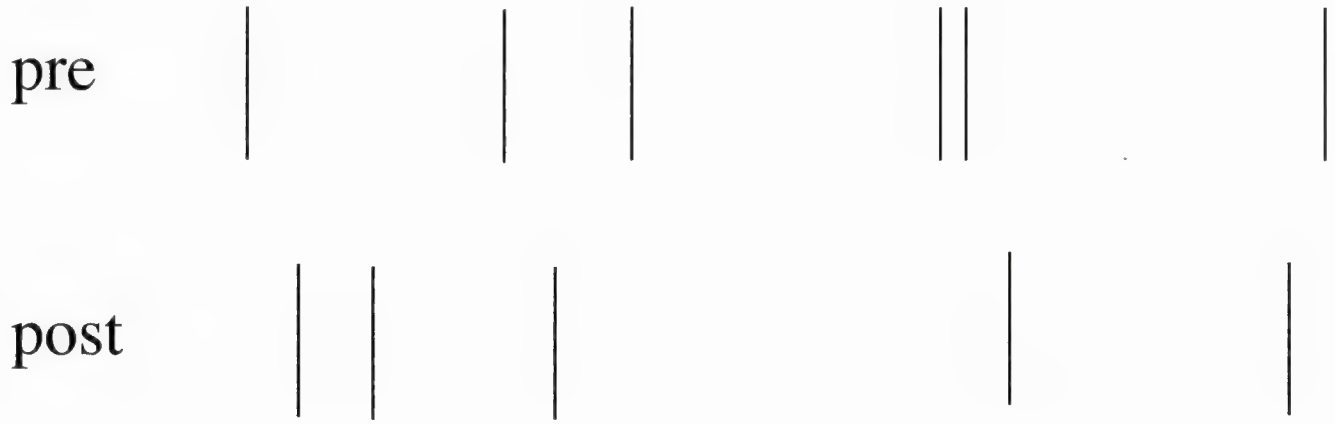
pre
post

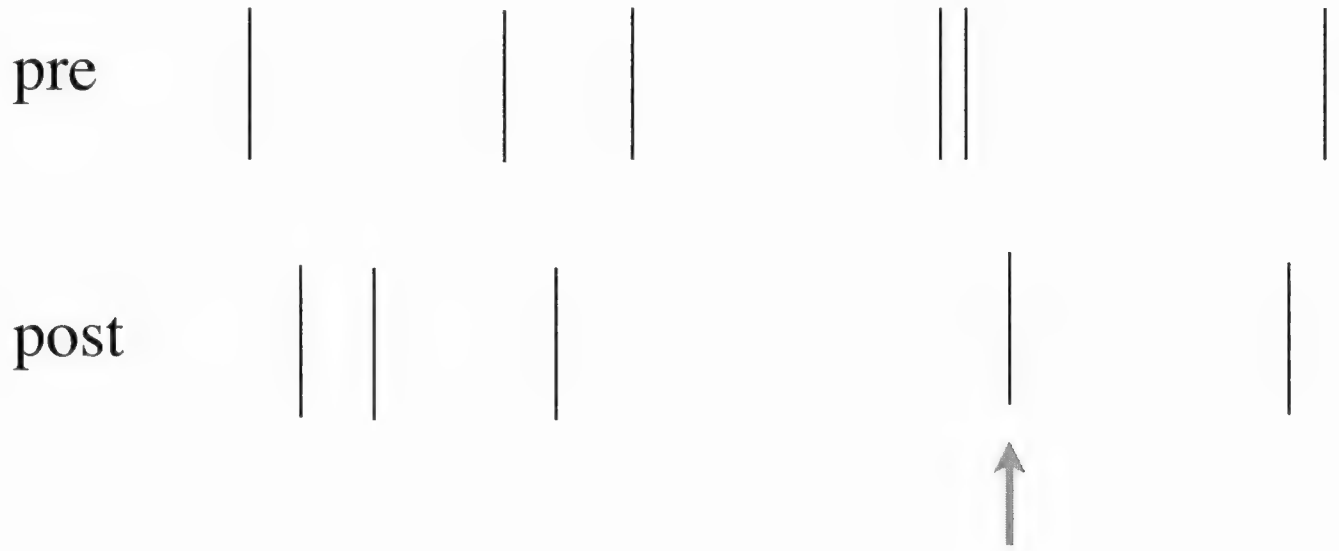
pre
post

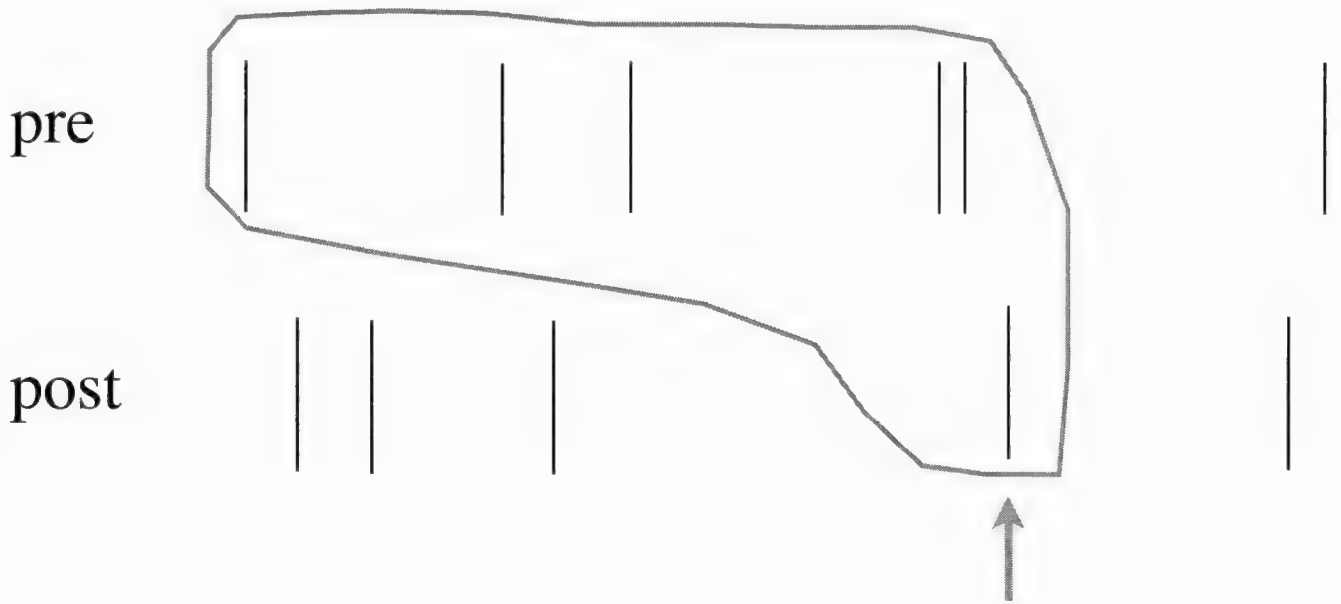
pre
post

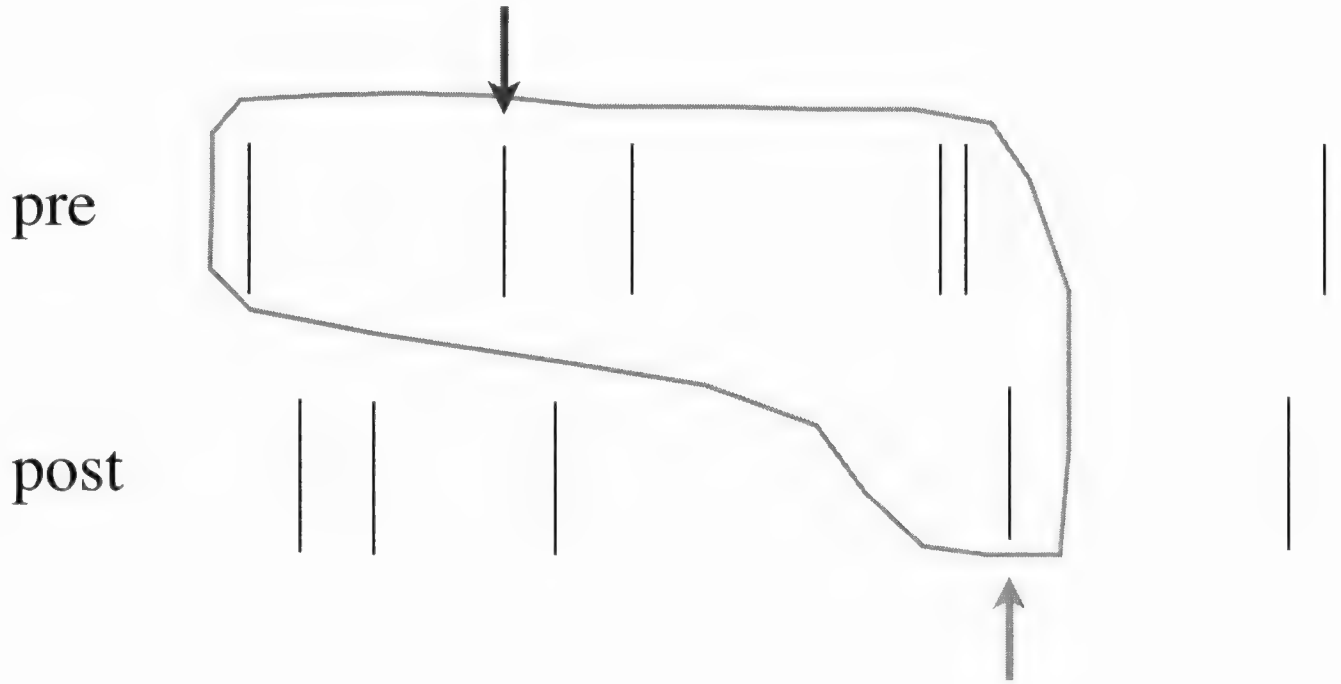
pre
post

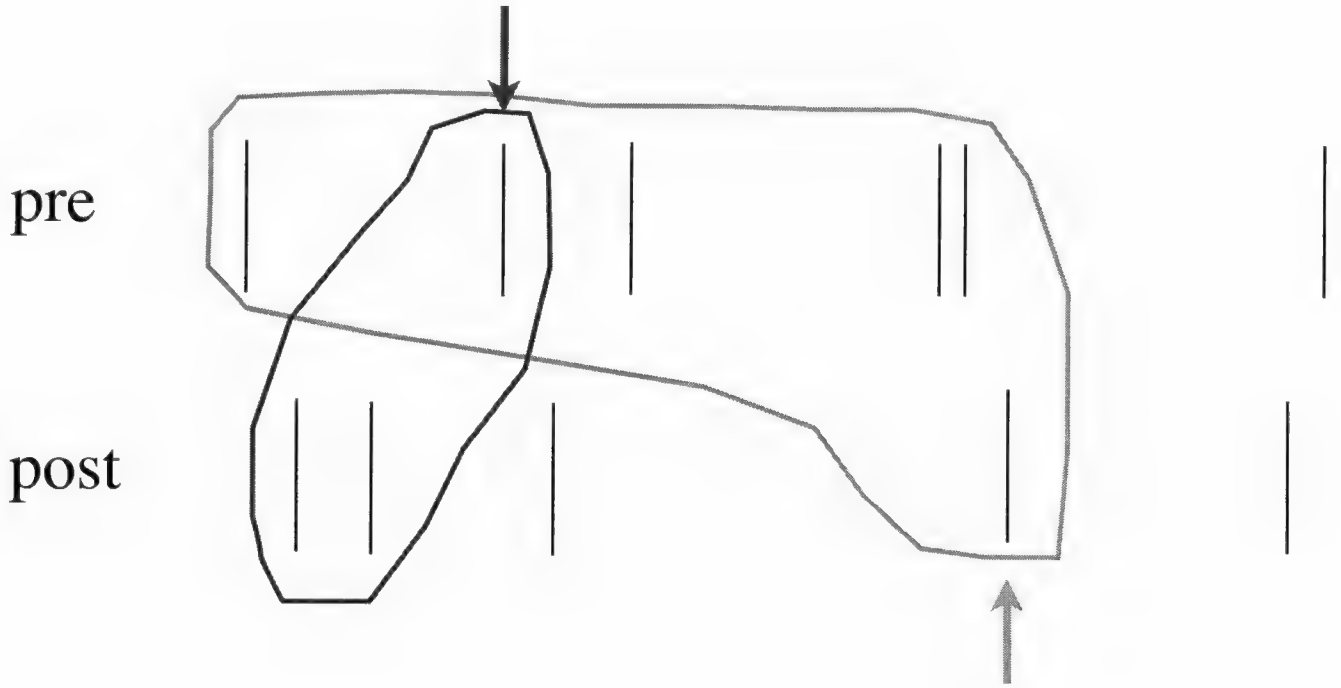
pre
post

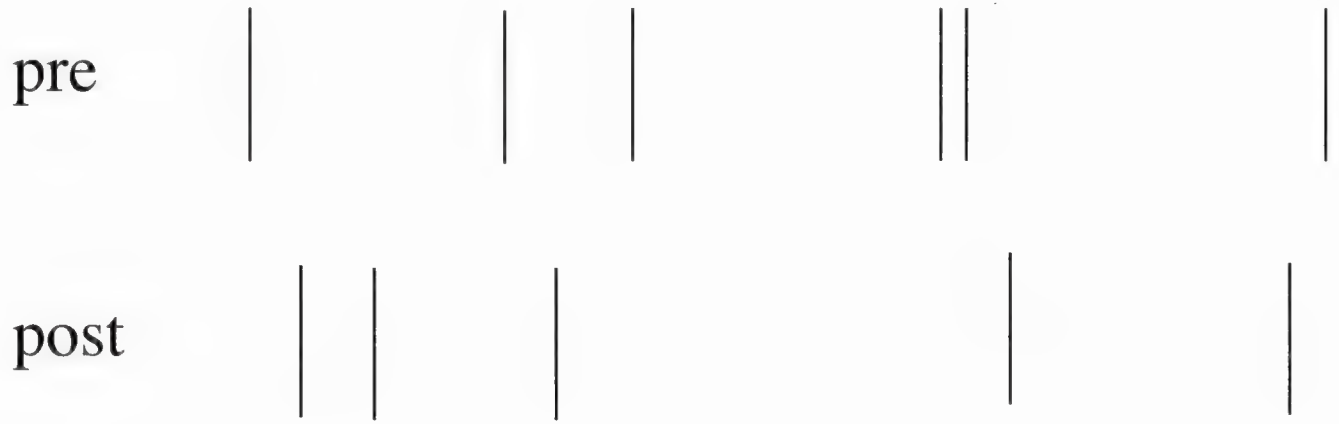

Sjostrom, Turrigiano, Nelson

Baktash Babadi

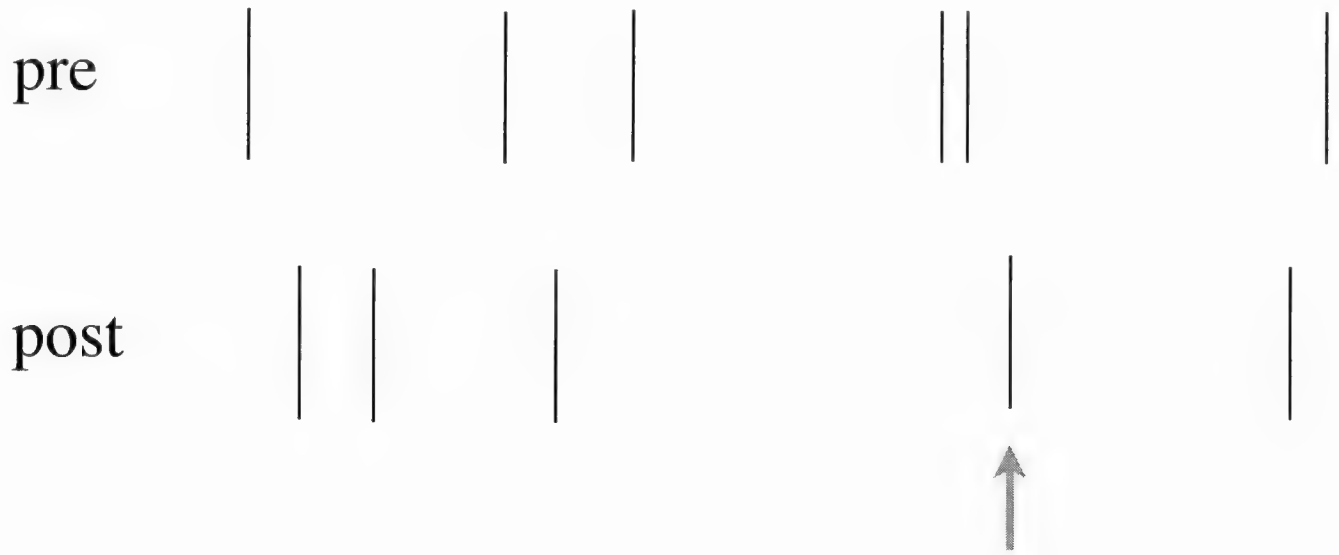

Sjostrom, Turrigiano, Nelson

Baktash Babadi

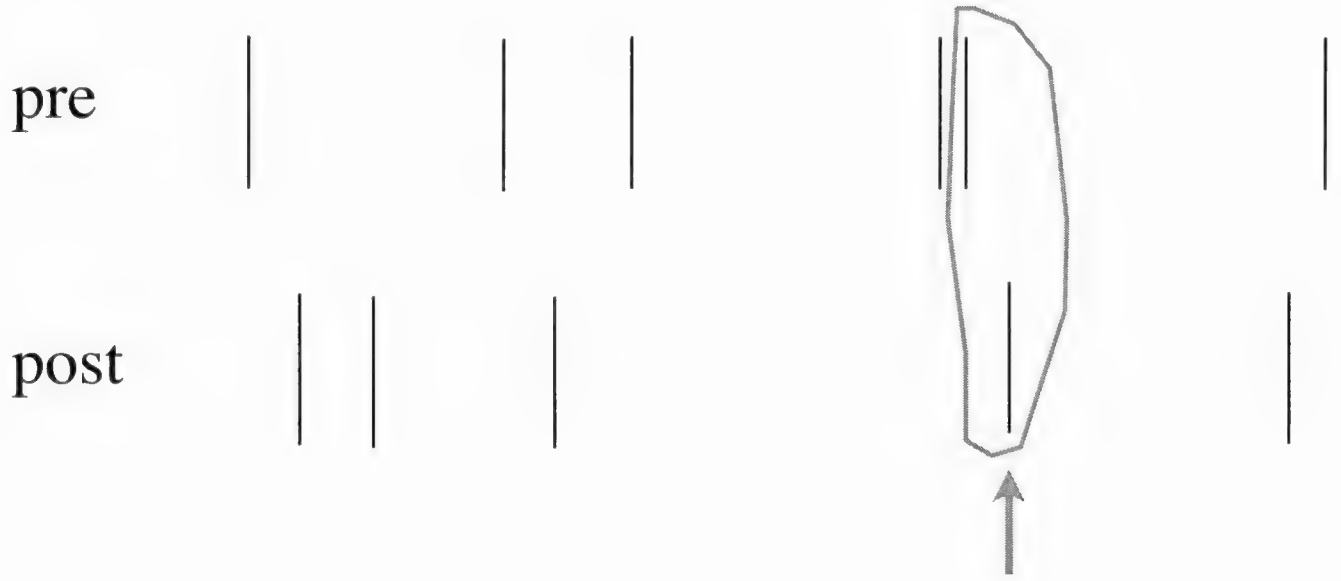

Sjostrom, Turrigiano, Nelson

Baktash Babadi

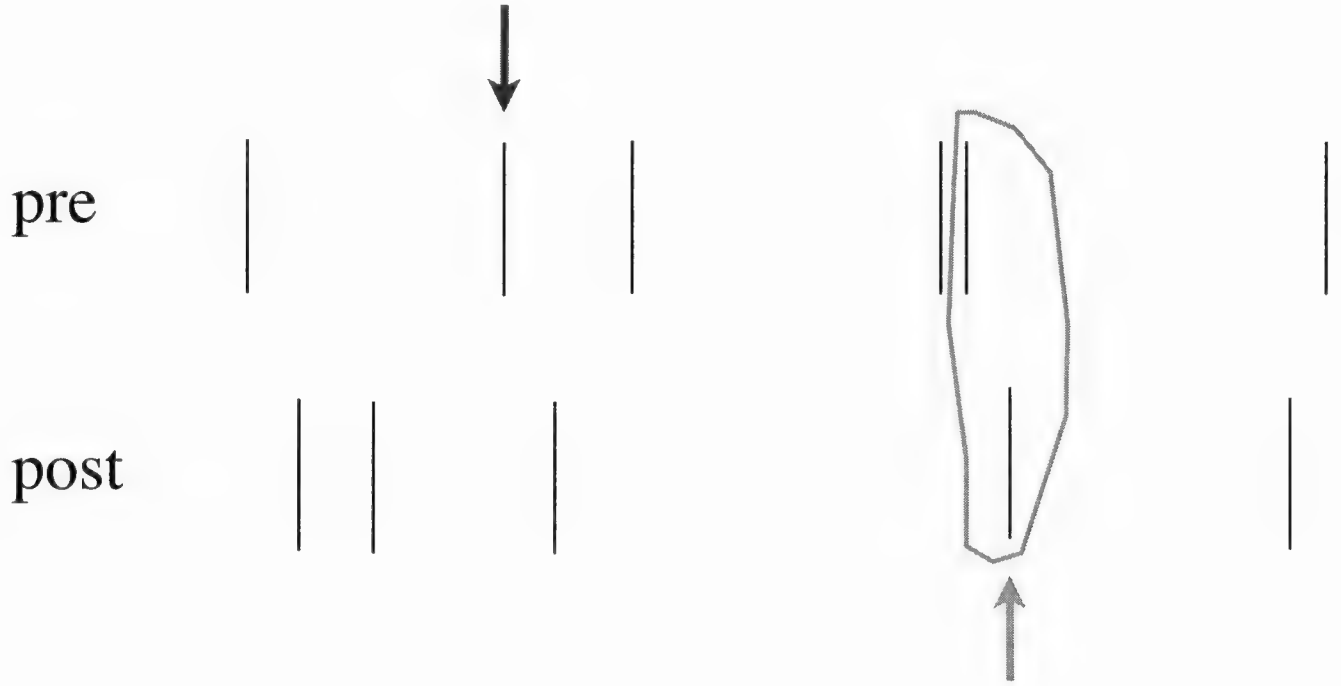

Sjostrom, Turrigiano, Nelson

Baktash Babadi

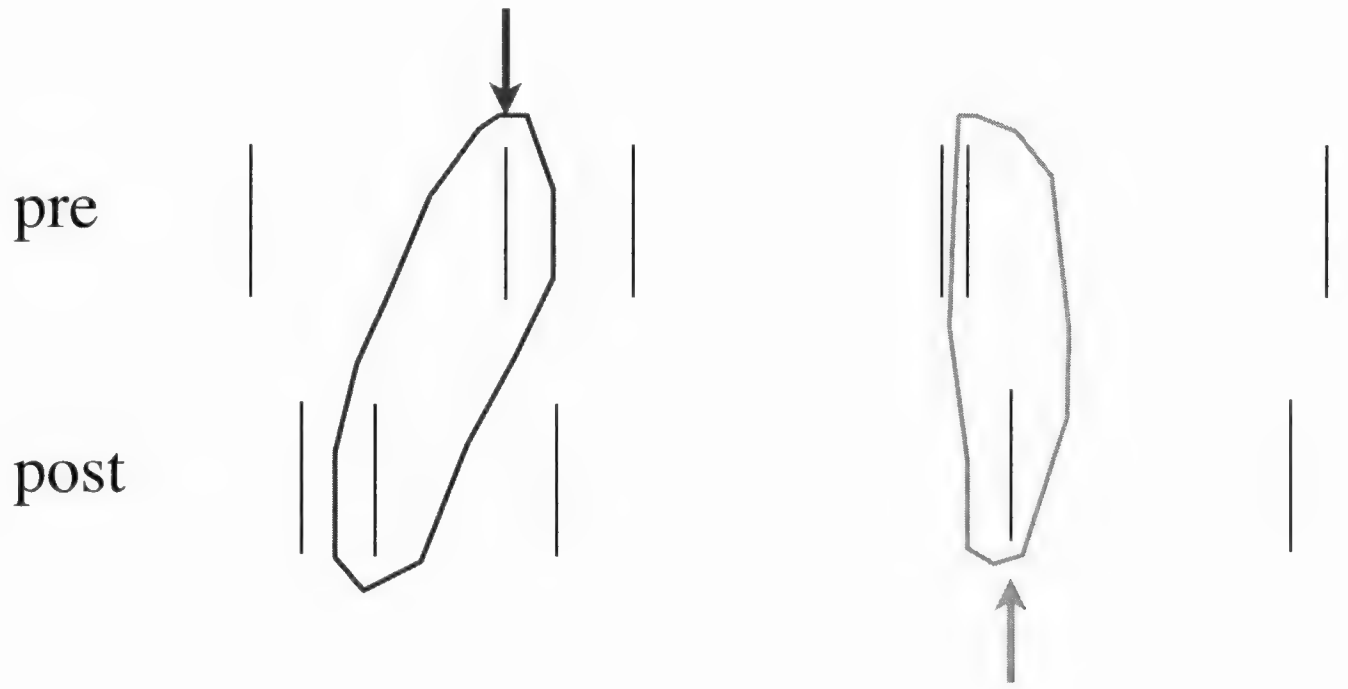

Sjostrom, Turrigiano, Nelson

Baktash Babadi

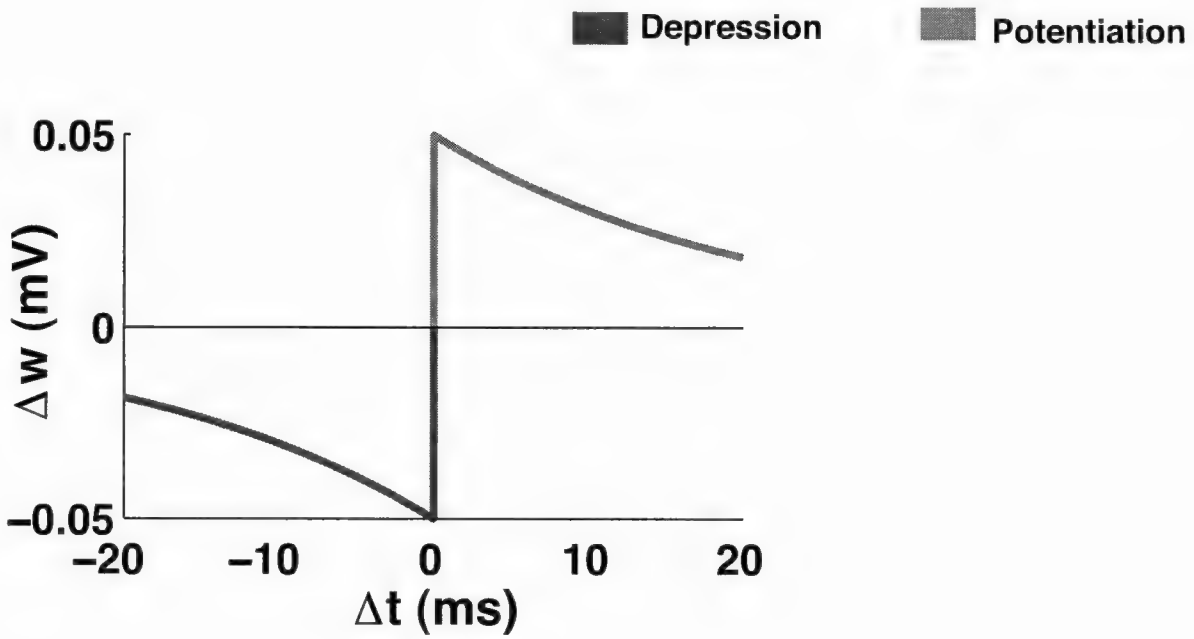
Depression
Potentiation
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)

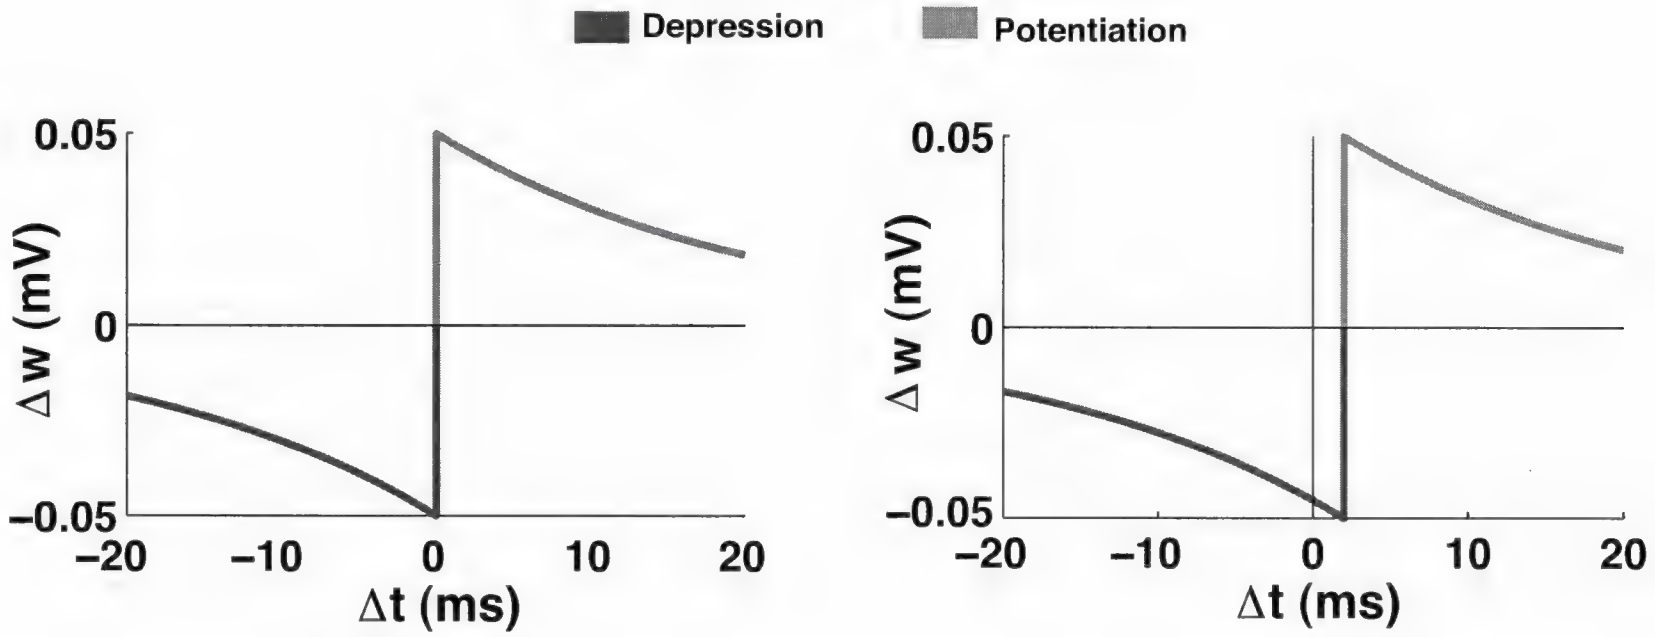
Depression
Potentiation
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)

$$\Delta w = F(\Delta t) = \begin{cases} -A_{\_}\, e^{(\Delta t - d)/\tau_{-}} & \text{if } \Delta t \leq d \\ \\ A_{+}\, e^{-(\Delta t - d)/\tau_{+}} & \text{if } \Delta t > d \end{cases}$$

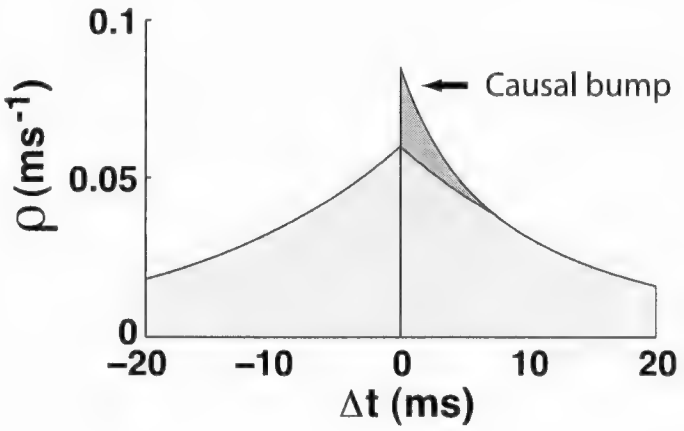
0.1
0.05
0
ρ (ms-1)
Causal bump
–20
–10
0
10
20
Δt (ms)

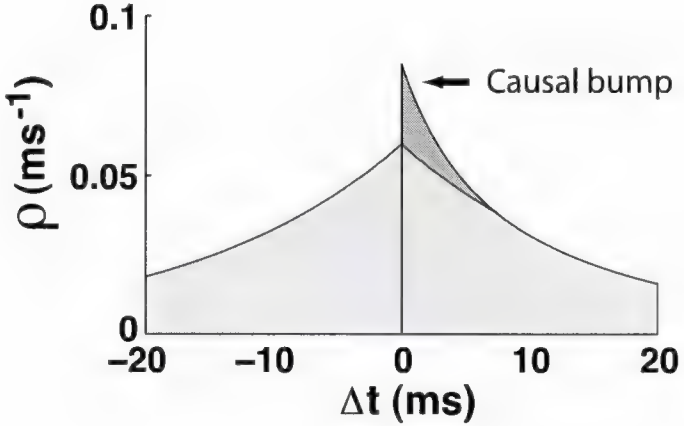
0.1
0.05
0
ρ (ms-1)
Causal bump
–20
–10
0
10
20
Δt (ms)

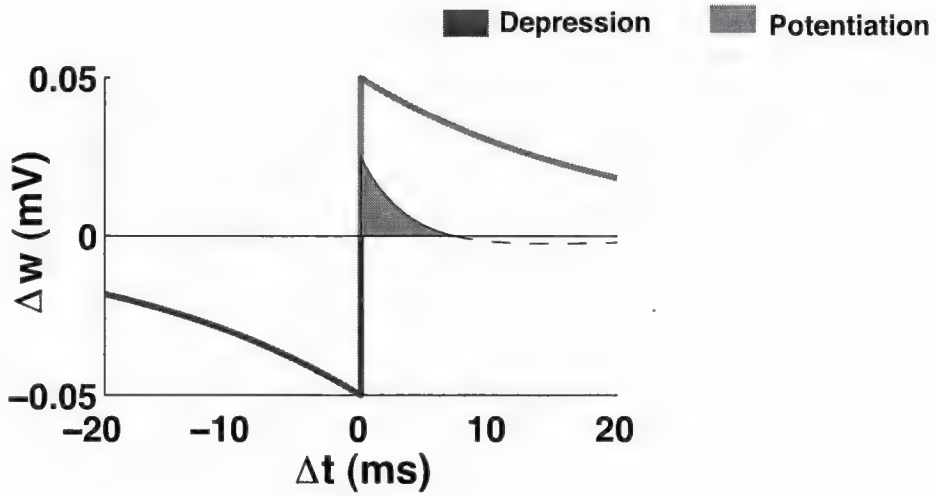
Depression
Potentiation
0.05
0
–0.05
Δw (mV)
–20
–10
0
10
20
Δt (ms)

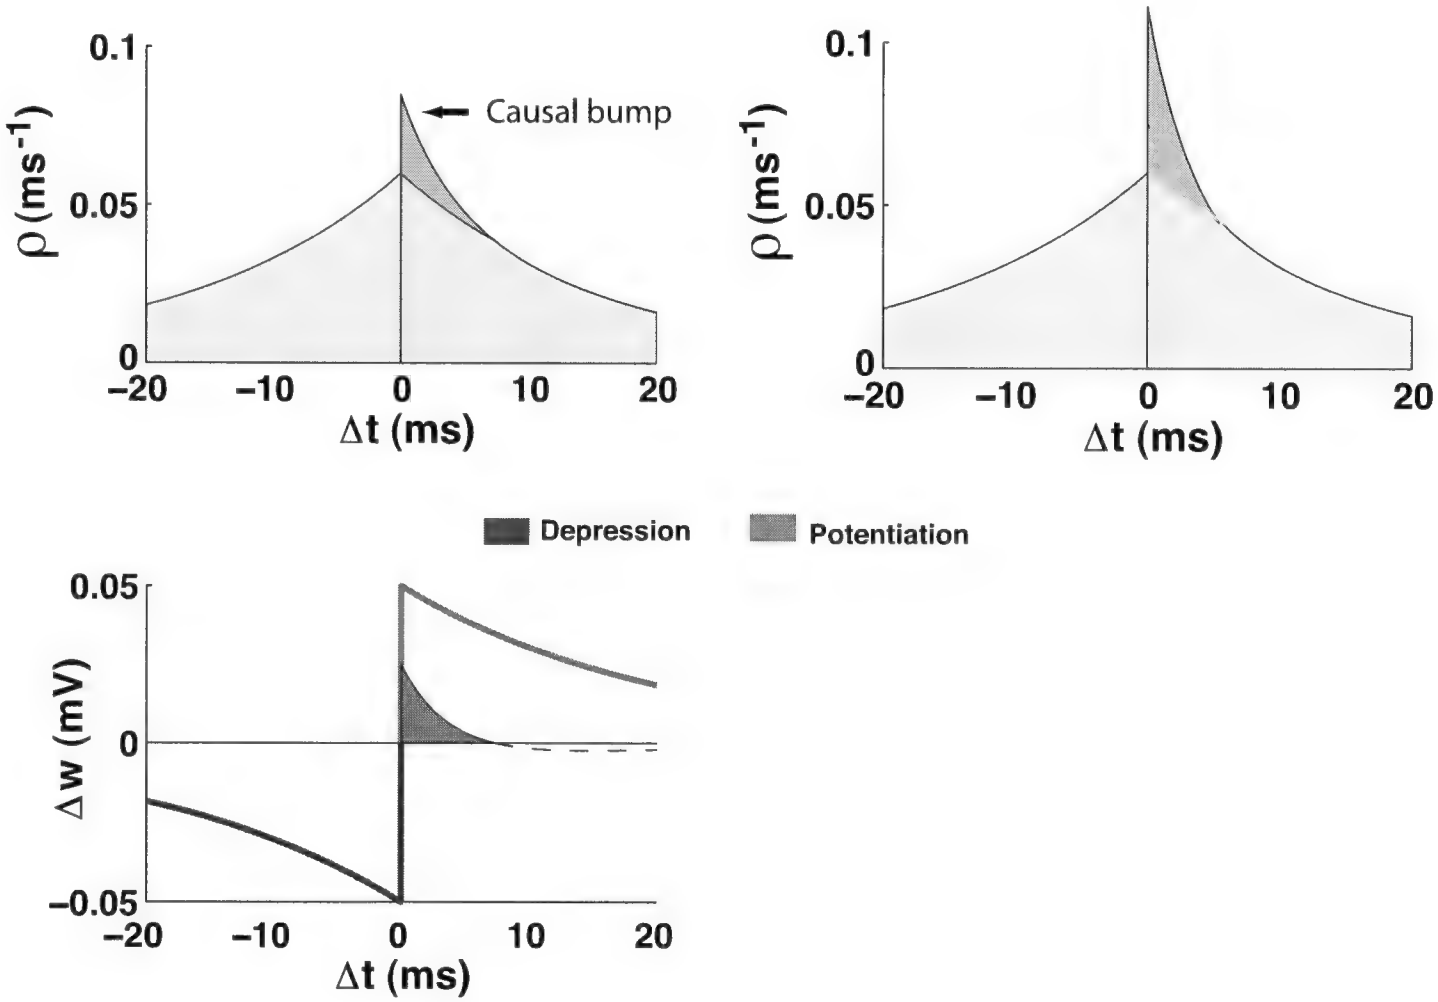
0.1
0.05
0
ρ (ms-1)
Causal bump
–20
–10
0
10
20
Δt (ms)
Depression
Potentiation
ΔW (mV)
–0.05

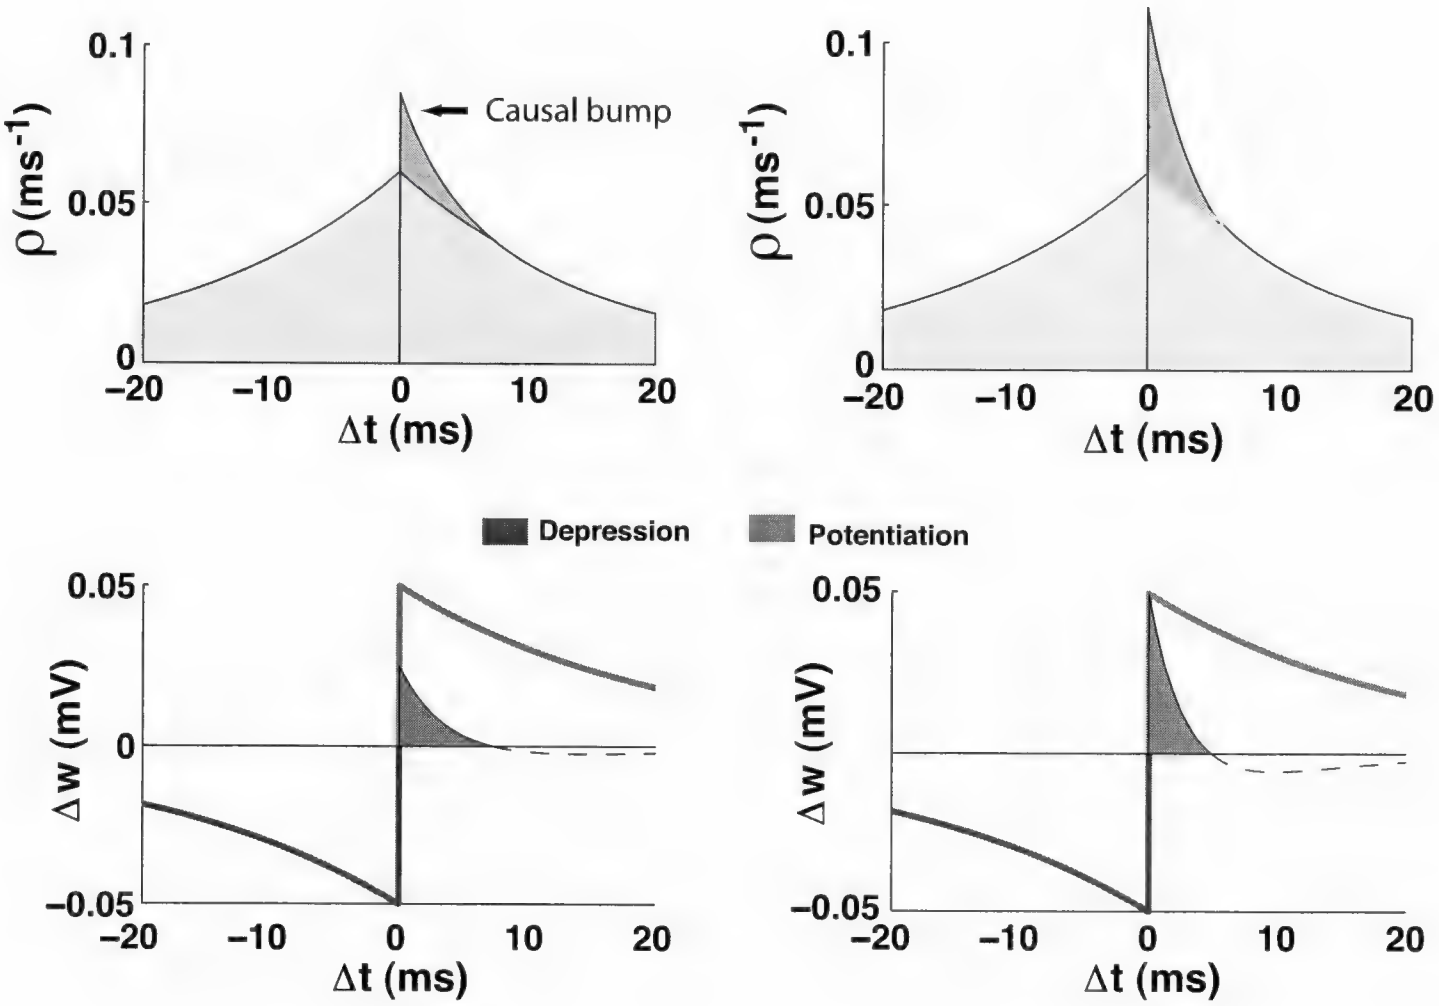
0.1
0.05
0
ρ (ms-1)
Causal bump
–20
–10
0
10
20
Δt (ms)
Depression
Potentiation
Δw (mV)
–0.05

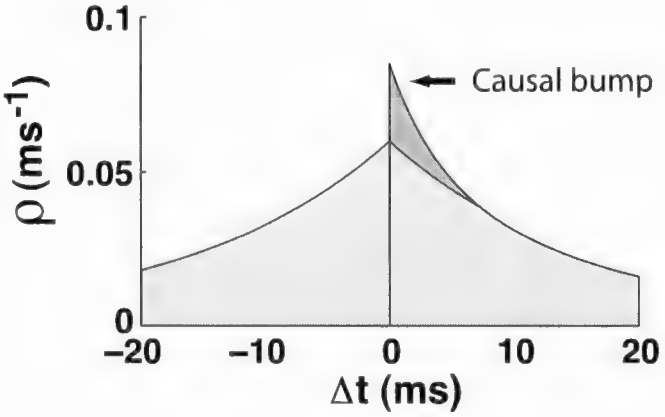
0.1
0.05
0
ρ (ms-1)
Causal bump
–20
–10
0
10
20
Δt (ms)

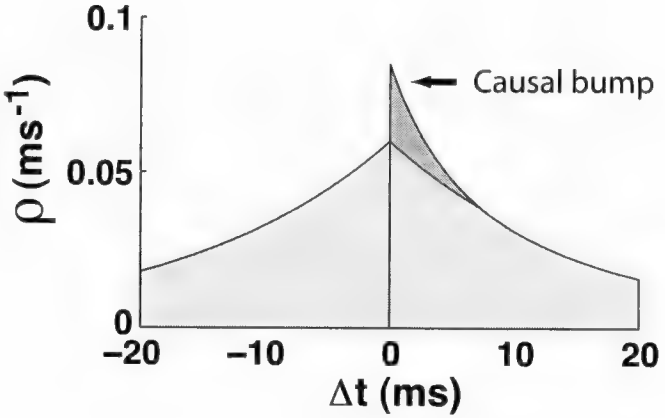
0.1
0.05
0
ρ (ms-1)
Causal bump
–20
–10
0
10
20
Δt (ms)

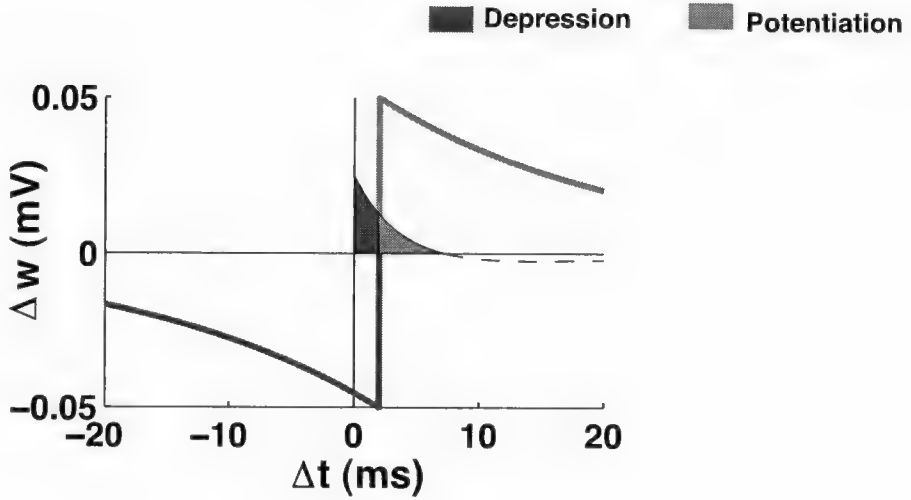
Depression
Potentiation
0.05
0
–0.05
Δ w (mV)
–20
–10
0
10
20
Δt (ms)

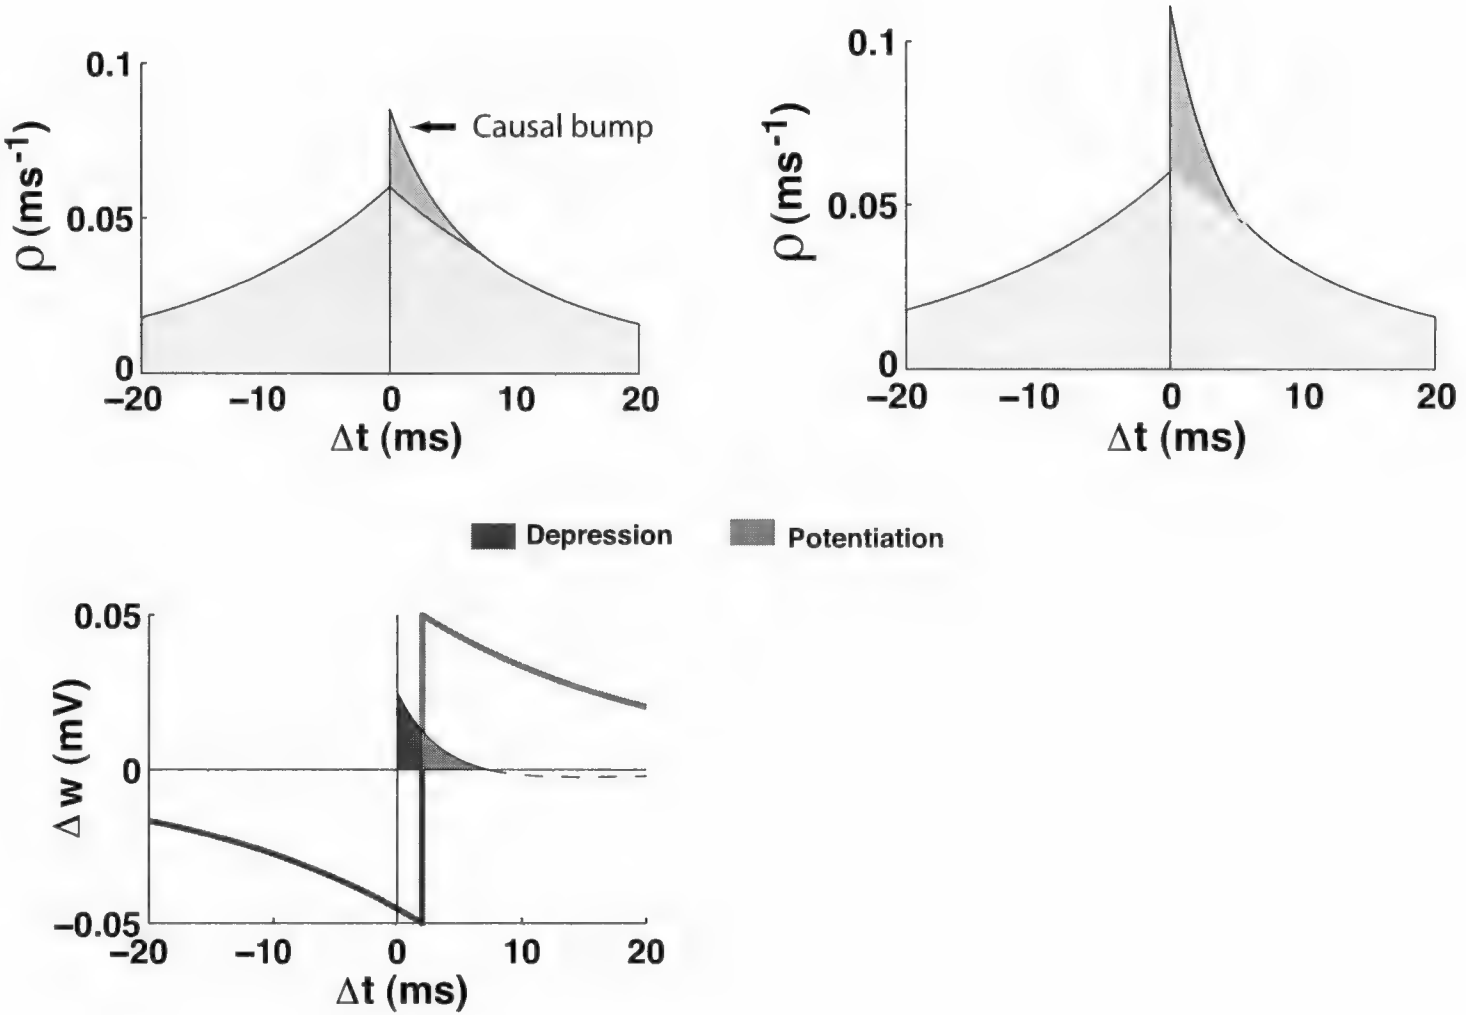
0.1
0.05
0
ρ (ms-1)
Causal bump
−20
−10
0
10
20
Δt (ms)
Depression
Potentiation
0.05
0
−0.05
Δ w (mV)

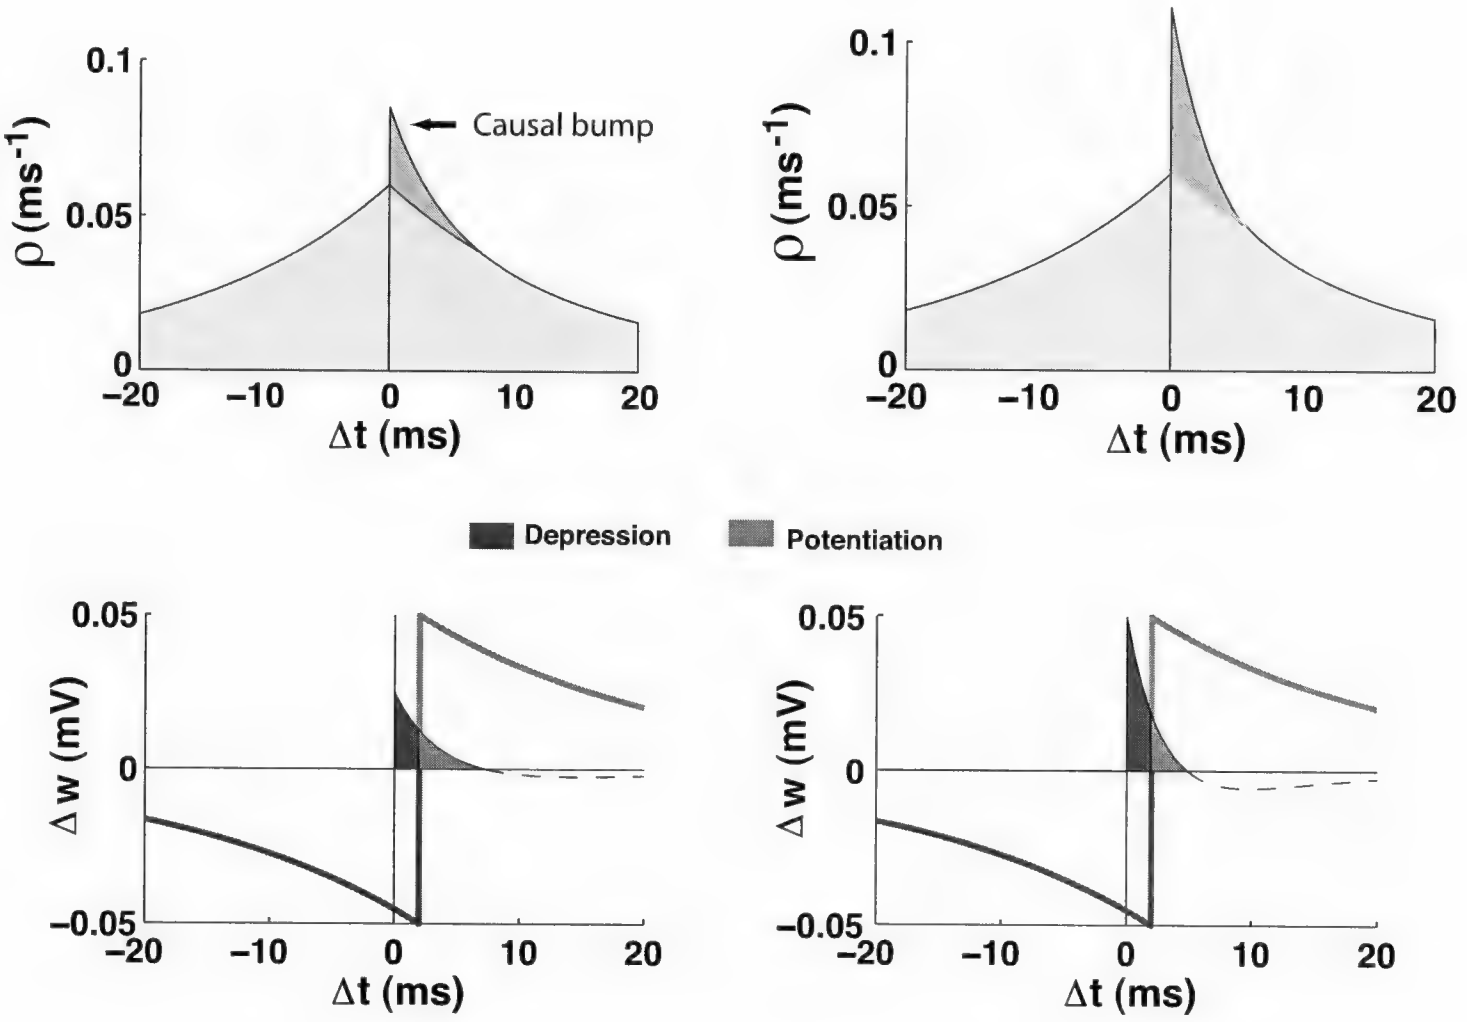
0.1
0.05
0
ρ (ms-1)
Causal bump
−20
−10
0
10
20
Δt (ms)
0.1
0.05
0
ρ (ms-1)
−20
−10
0
10
20
Δt (ms)
Depression
Potentiation
0.05
0
−0.05
Δ w (mV)
−20
−10
0
10
20
Δt (ms)
0.05
0
−0.05
Δ w (mV)
−20
−10
0
10
20
Δt (ms)

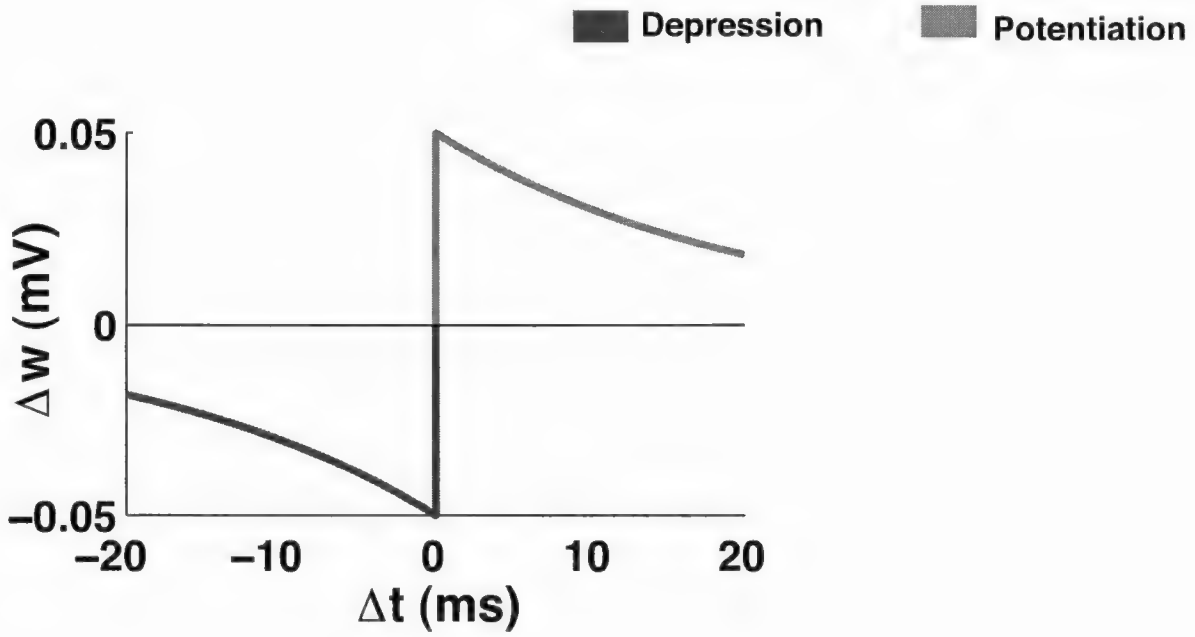
Depression
Potentiation
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)

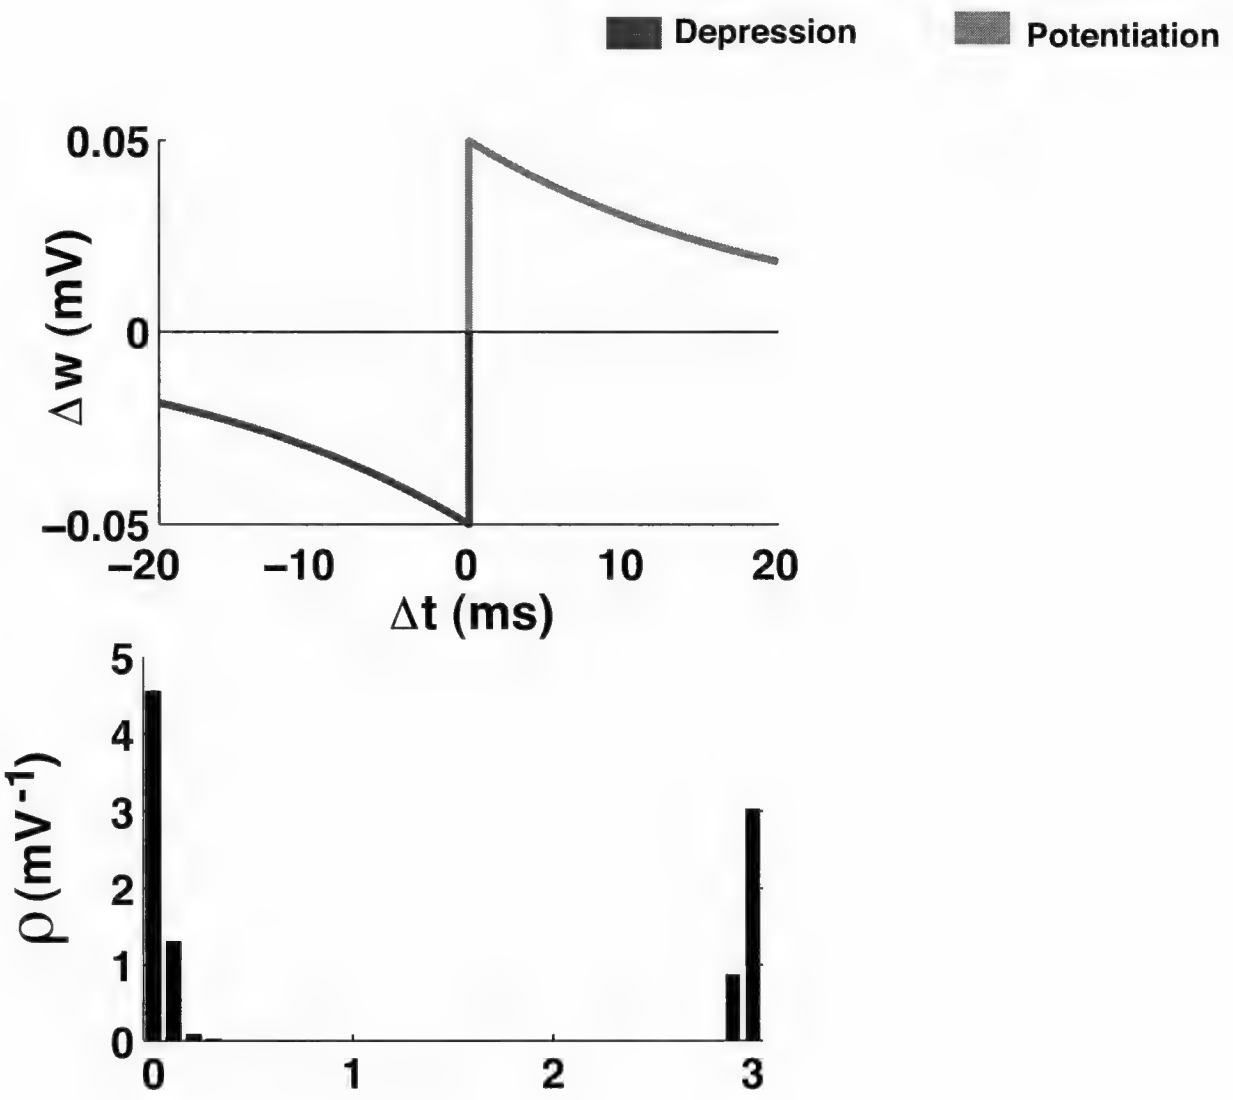
Depression
Potentiation
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)
5
4
3
2
1
0
ρ (mV-1)
0
1
2
3

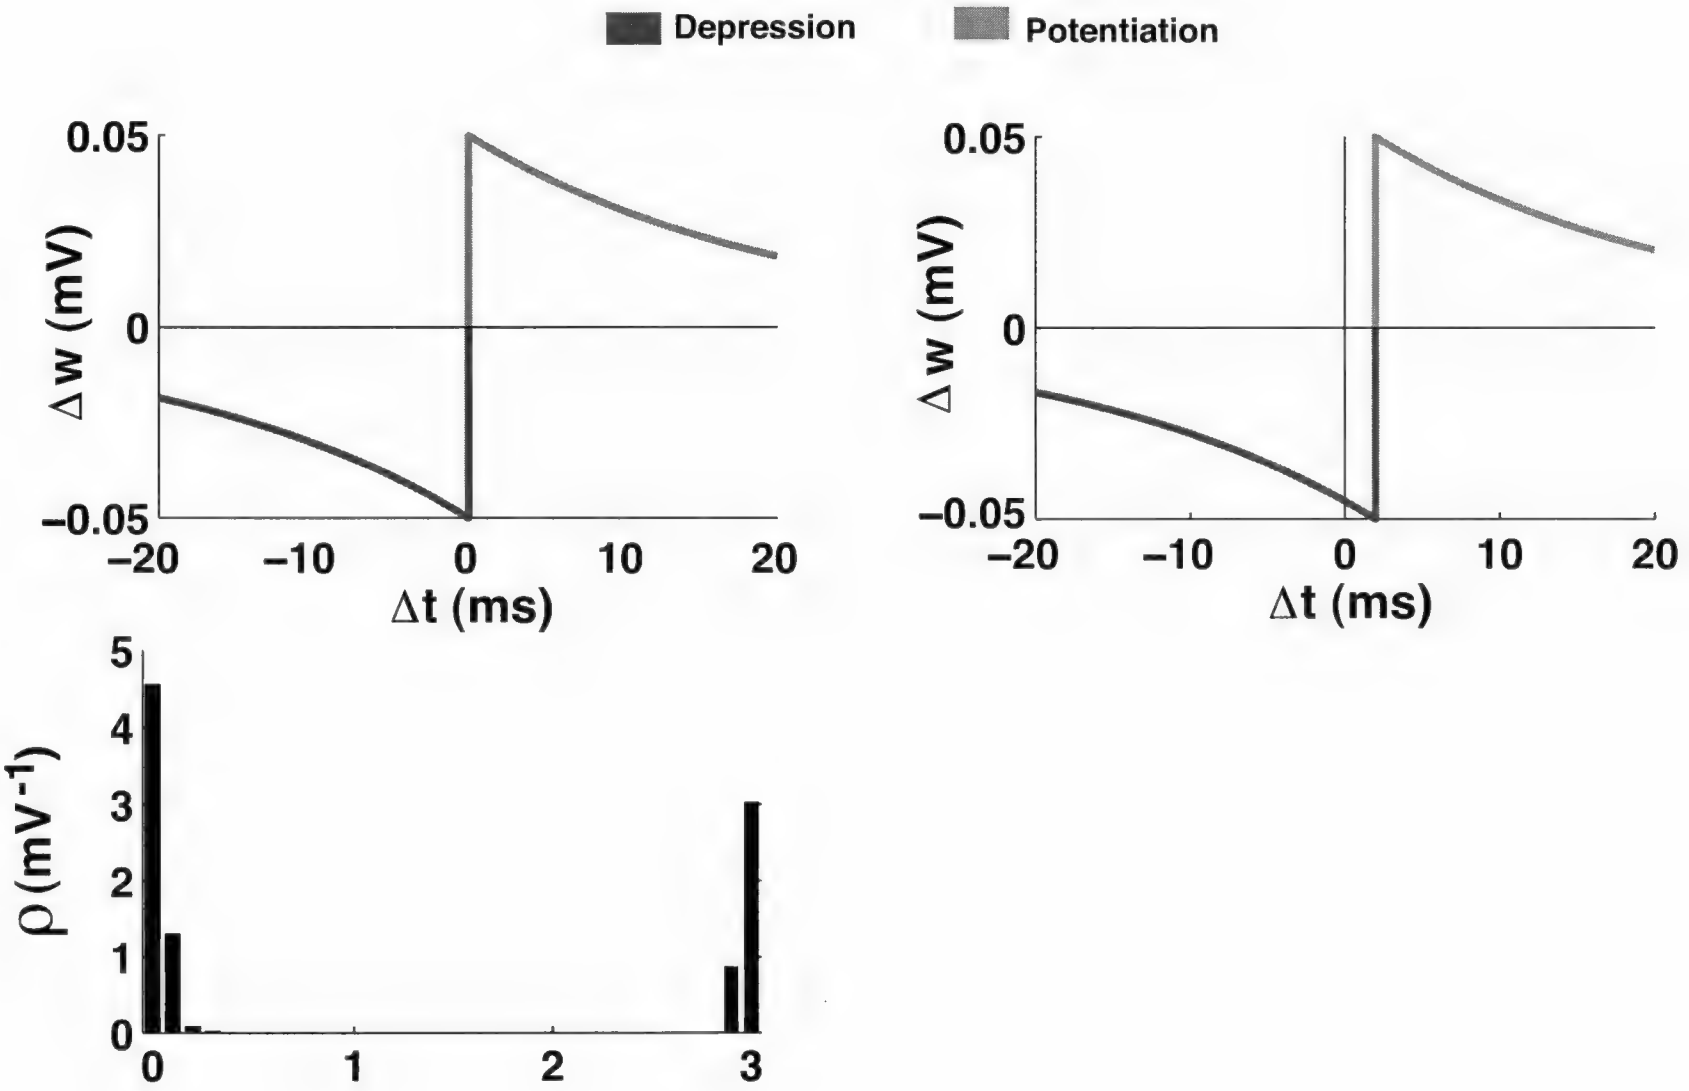
Depression
Potentiation
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)
ρ (mV-1)
5
4
3
2
1
0
0
1
2
3

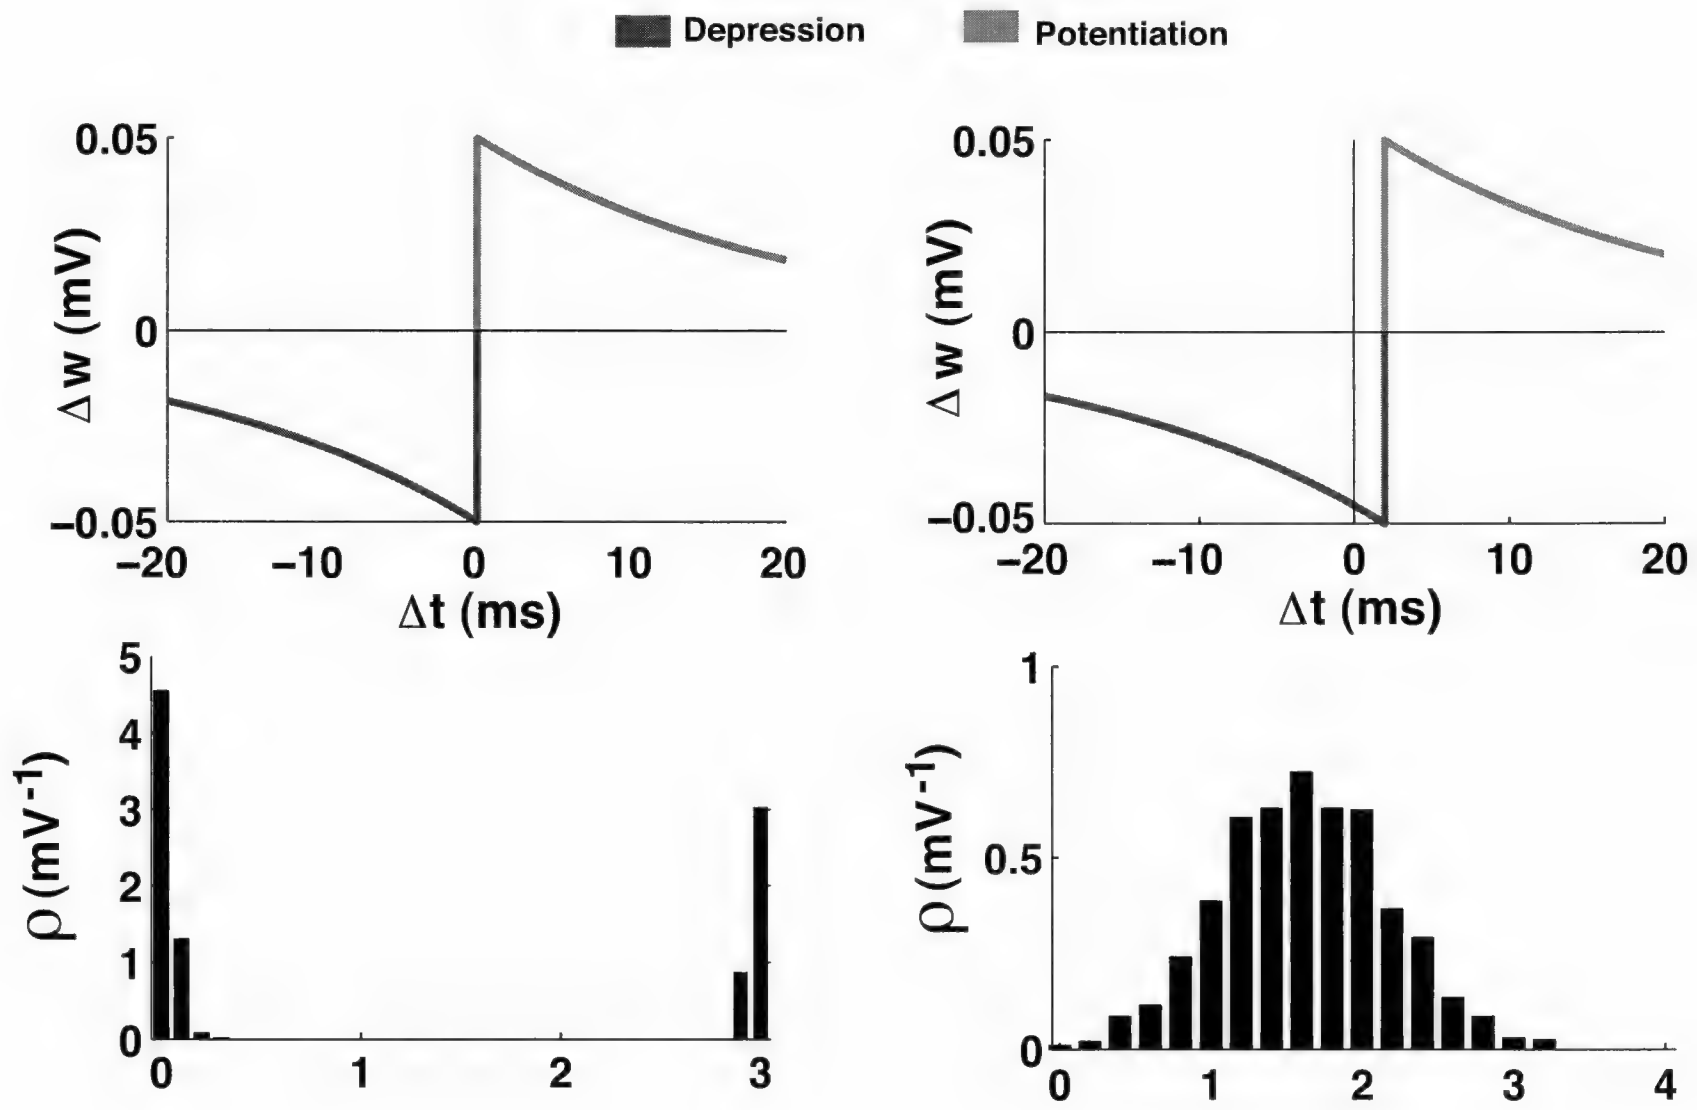
Depression
Potentiation
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)
0.05
0
−0.05
Δw (mV)
−20
−10
0
10
20
Δt (ms)
ρ (mV-1)
5
4
3
2
1
0
0
1
2
3
ρ (mV-1)
1
0.5
0
0
1
2
3
4

Shifted-STDP Demo

$$\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial F}{\partial w}$$

$$\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial F}{\partial w} \qquad F = D_1(w)\rho(w,t) - \frac{1}{2}\frac{\partial D_2(w)\rho(w,t)}{\partial w}$$

$$\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial F}{\partial w} \qquad F = D_1(w)\rho(w,t) - \frac{1}{2}\frac{\partial D_2(w)\rho(w,t)}{\partial w}$$

$$\boxed{\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial}{\partial w}\big(D_1(w)\rho(w,t)\big) + \frac{1}{2}\frac{\partial^2}{\partial w^2}\big(D_2(w)\rho(w,t)\big)}$$

$$\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial F}{\partial w} \qquad F = D_1(w)\rho(w,t) - \frac{1}{2}\frac{\partial D_2(w)\rho(w,t)}{\partial w}$$

$$\boxed{\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial}{\partial w}\big(D_1(w)\rho(w,t)\big) + \frac{1}{2}\frac{\partial^2}{\partial w^2}\big(D_2(w)\rho(w,t)\big)}$$

$$\boxed{\rho(w) = \frac{N_1}{D_2(w)} \exp\left(\int_0^{\omega} d\omega' \frac{2D_1(w')}{D_2(w')}\right)}$$

$$\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial F}{\partial w} \qquad F = D_1(w)\rho(w,t) - \frac{1}{2}\frac{\partial D_2(w)\rho(w,t)}{\partial w}$$

$$\boxed{\frac{\partial \rho(w,t)}{\partial t} = -\frac{\partial}{\partial w}\big(D_1(w)\rho(w,t)\big) + \frac{1}{2}\frac{\partial^2}{\partial w^2}\big(D_2(w)\rho(w,t)\big)}$$

$$\boxed{\rho(w) = \frac{N_1}{D_2(w)} \exp\left(\int_0^{\omega} d\omega' \frac{2D_1(w')}{D_2(w')}\right)}$$

$$D_1 = \int_{-\infty}^{\infty} dt\ P(t)F(t) \quad \text{and} \quad D_2 = \int_{-\infty}^{\infty} dt\ P(t)\,F^2(t)$$

$$P(t) = \left(r_{\text{ex}} + r_{\text{post}}(t)\right) \exp\left(-\int_0^t ds \ \left(r_{\text{ex}} + r_{\text{post}}(s)\right)\right)$$

$$P(t) = \left(r_{\text{ex}} + r_{\text{post}}(t)\right) \exp\left(-\int_0^t ds \, \left(r_{\text{ex}} + r_{\text{post}}(s)\right)\right)$$

$$r_{\text{post}}(t) \approx \bar{r}_{\text{post}} + \frac{w \exp(-t/\tau_{\text{s}})}{(V_{\text{th}} - V_r)\,\tau_{\text{m}}}$$

$$P(t) = \left(r_{\text{ex}} + r_{\text{post}}(t)\right) \exp\left(-\int_0^t ds \ \left(r_{\text{ex}} + r_{\text{post}}(s)\right)\right)$$

$$r_{\text{post}}(t) \approx \bar{r}_{\text{post}} + \frac{w \exp(-t/\tau_{\text{s}})}{(V_{\text{th}} - V_r)\,\tau_{\text{m}}}$$

$$\bar{r}_{\text{post}} \approx \frac{N_{\text{ex}}\, r_{\text{ex}}\, \tau_{\text{s}}\, \langle w \rangle - N_{\text{in}}\, r_{\text{in}}\, \tau_{\text{s}}\, w_{\text{in}} - 0.5\,(V_{\text{th}} - V_r)}{(V_{\text{th}} - V_r)\,\tau_{\text{m}}}$$

$$P(t) = \left(r_{\text{ex}} + r_{\text{post}}(t)\right) \exp\left(-\int_0^t ds \, \left(r_{\text{ex}} + r_{\text{post}}(s)\right)\right)$$

$$r_{\text{post}}(t) \approx \bar{r}_{\text{post}} + \frac{w \exp(-t/\tau_{\text{s}})}{(V_{\text{th}} - V_r)\,\tau_{\text{m}}}$$

$$\bar{r}_{\text{post}} \approx \frac{N_{\text{ex}}\, r_{\text{ex}}\, \tau_{\text{s}}\, \langle w \rangle - N_{\text{in}}\, r_{\text{in}}\, \tau_{\text{s}}\, w_{\text{in}} - 0.5\,(V_{\text{th}} - V_r)}{(V_{\text{th}} - V_r)\,\tau_{\text{m}}}$$

$$P(t) \approx \begin{cases} r \exp(r\,t) & \text{if } t \leq 0 \\ r \exp(-r\,t) + w \dfrac{\tau_{\text{s}}}{\tau_{\text{m}}(V_{\text{th}} - V_r)} \left(\lambda \exp(-\lambda\,t) - r \exp(-r\,t)\right) & \text{if } t > 0 \end{cases}$$

with

$$\lambda = r + \frac{1}{\tau_{\text{s}}}\,.$$

$$\rho(w) = N_0(w + \mu)^{k-1} \exp\left(-\frac{w + \mu}{\theta}\right)$$

$$\boxed{\rho(w) = N_0(w+\mu)^{k-1}\exp\left(-\frac{w+\mu}{\theta}\right)}$$

$$\mu = \frac{\delta}{\gamma}, \quad k = \frac{2(\beta\gamma - \alpha\delta)}{\gamma^2} \quad \text{and} \quad \theta = -\frac{\gamma}{2\alpha}$$

$$\begin{aligned}
\alpha &= \frac{1}{\tau_m(V_{th}-V_r)}\left(\frac{A_+(\tau_+\tau_s + d\,\tau_s)}{(1+r\tau_+)(r\tau_s\tau_+ + \tau_s + \tau_+)} - (A_+ + A_-)d\right) \\
\beta &= A_+ r\tau_+\left(\frac{1-dr}{1+r\,\tau_+}\right) - A_- r\,\tau_-\left(\frac{1+d\,r}{1+r\,\tau_-}\right) \\
\gamma &= \frac{1}{\tau_m(V_{th}-V_r)}\left(\frac{A_+^2(2\tau_+\tau_s + 4d\,\tau_s)}{(2+r\tau_+)(r\tau_s\tau_+ + 2\tau_s + \tau_+)} - (A_+^2 - A_-^2)d\right) \\
\delta &= A_+^2 r\tau_+\left(\frac{1-d\,r}{2+r\tau_+}\right) + A_-^2 r\tau_-\left(\frac{1+dr}{2+r\tau_-}\right).
\end{aligned}$$

$$\rho(w) = N_0(w+\mu)^{k-1}\exp\left(-\frac{w+\mu}{\theta}\right)$$

$$\mu = \frac{\delta}{\gamma}, \quad k = \frac{2(\beta\gamma - \alpha\delta)}{\gamma^2} \quad \text{and} \quad \theta = -\frac{\gamma}{2\alpha}$$

$$\begin{aligned}
\alpha &= \frac{1}{\tau_m(V_{th}-V_r)}\left(\frac{A_+(\tau_+\tau_s + d\,\tau_s)}{(1+r\tau_+)(r\tau_s\tau_+ + \tau_s + \tau_+)} - (A_+ + A_-)d\right) \\
\beta &= A_+ r\tau_+\left(\frac{1-dr}{1+r\,\tau_+}\right) - A_- r\,\tau_-\left(\frac{1+d\,r}{1+r\,\tau_-}\right) \\
\gamma &= \frac{1}{\tau_m(V_{th}-V_r)}\left(\frac{A_+^2(2\tau_+\tau_s + 4d\,\tau_s)}{(2+r\tau_+)(r\tau_s\tau_+ + 2\tau_s + \tau_+)} - (A_+^2 - A_-^2)d\right) \\
\delta &= A_+^2 r\tau_+\left(\frac{1-d\,r}{2+r\tau_+}\right) + A_-^2 r\tau_-\left(\frac{1+dr}{2+r\tau_-}\right).
\end{aligned}$$

$\alpha < 0$

$\beta > 0$

$\gamma > 0$

$\delta > 0$

$$\boxed{\rho(w) = N_0 (w+\mu)^{k-1} \exp\left(-\frac{w+\mu}{\theta}\right)}$$

$$\mu = \frac{\delta}{\gamma}, \quad k = \frac{2(\beta\gamma - \alpha\delta)}{\gamma^2} \quad \text{and} \quad \theta = -\frac{\gamma}{2\alpha}$$

$$\begin{aligned}
\alpha &= \frac{1}{\tau_m (V_{th} - V_r)} \left(\frac{A_+(\tau_+\tau_s + d\tau_s)}{(1+r\tau_+)(r\tau_s\tau_+ + \tau_s + \tau_+)} - (A_+ + A_-)d \right) \\
\beta &= A_+ r\tau_+ \left(\frac{1-dr}{1+r\tau_+} \right) - A_- r\tau_- \left(\frac{1+dr}{1+r\tau_-} \right) \\
\gamma &= \frac{1}{\tau_m (V_{th} - V_r)} \left(\frac{A_+^2(2\tau_+\tau_s + 4d\tau_s)}{(2+r\tau_+)(r\tau_s\tau_+ + 2\tau_s + \tau_+)} - (A_+^2 - A_-^2)d \right) \\
\delta &= A_+^2 r\tau_+ \left(\frac{1-dr}{2+r\tau_+} \right) + A_-^2 r\tau_- \left(\frac{1+dr}{2+r\tau_-} \right).
\end{aligned}$$

$$\boxed{\begin{array}{c} \alpha < 0 \\ \beta > 0 \\ \gamma > 0 \\ \delta > 0 \end{array}}$$

$$\boxed{d > 0 \qquad A_+\tau_+ > A_-\tau_-}$$

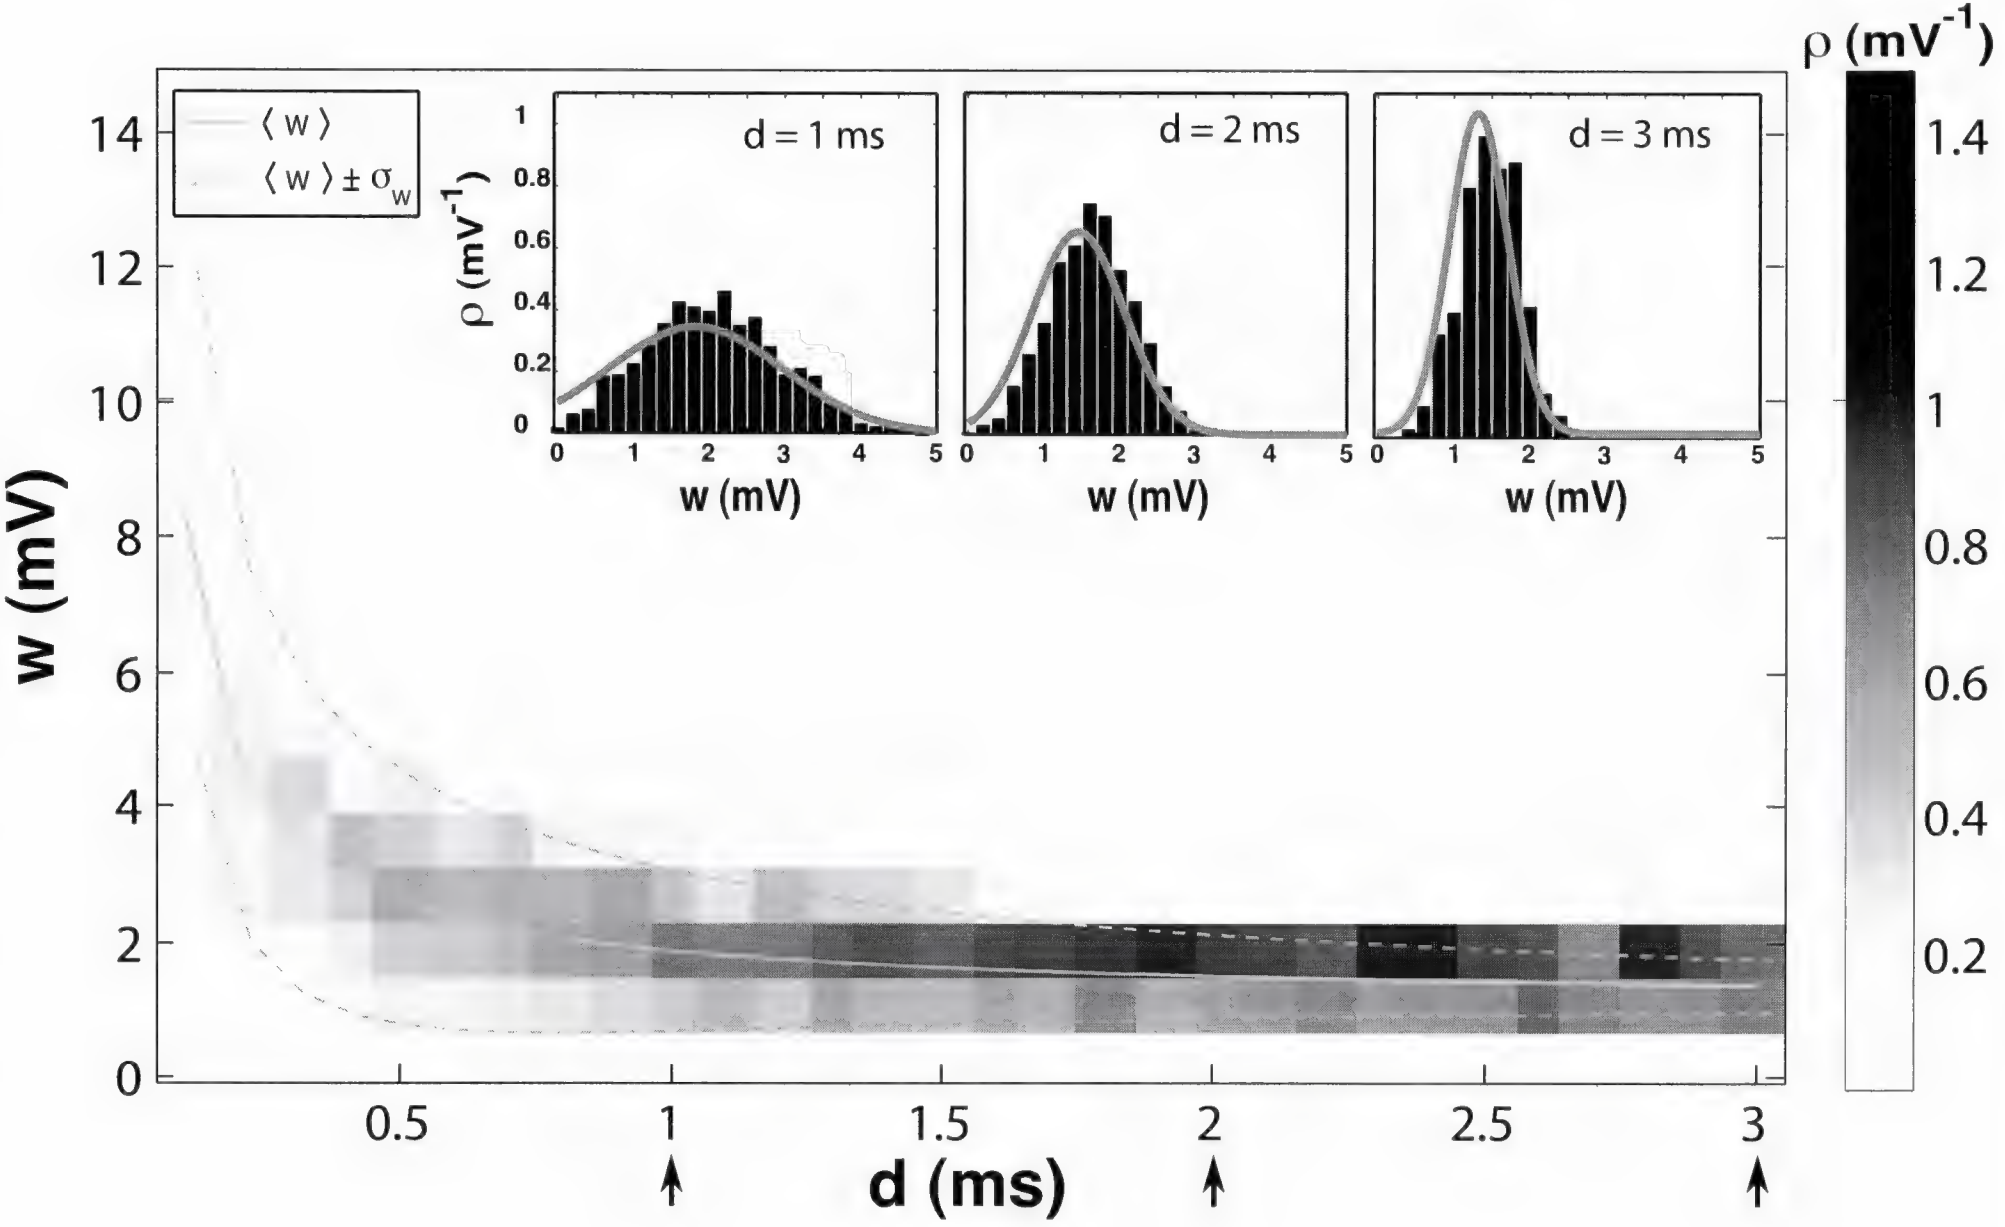
⟨ w ⟩
⟨ w ⟩ ± σ_w
d = 1 ms
d = 2 ms
d = 3 ms
ρ (mV^-1)
w (mV)
d (ms)
ρ (mV^-1)

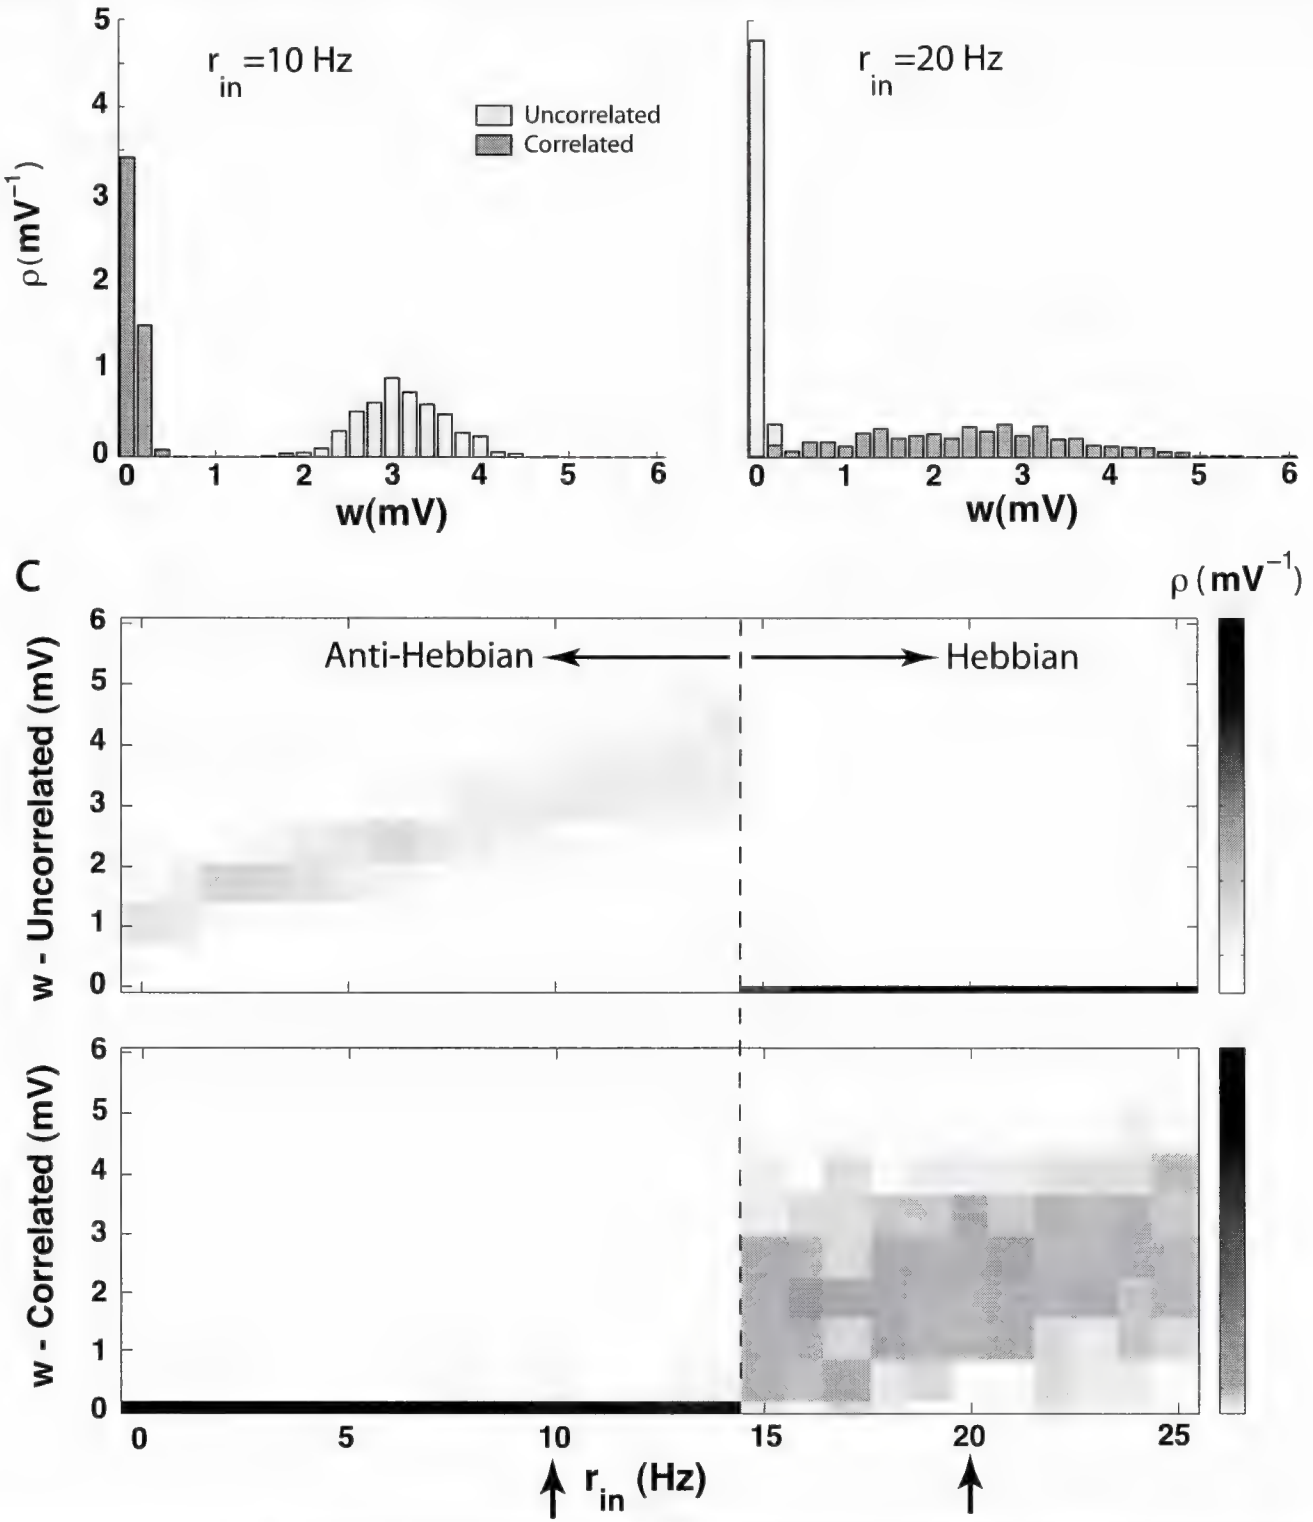
r in =10 Hz
Uncorrelated
Correlated
ρ(mV⁻¹)
w(mV)
r in =20 Hz
w(mV)
C
ρ (mV⁻¹)
Anti-Hebbian
Hebbian
w - Uncorrelated (mV)
w - Correlated (mV)
r in (Hz)

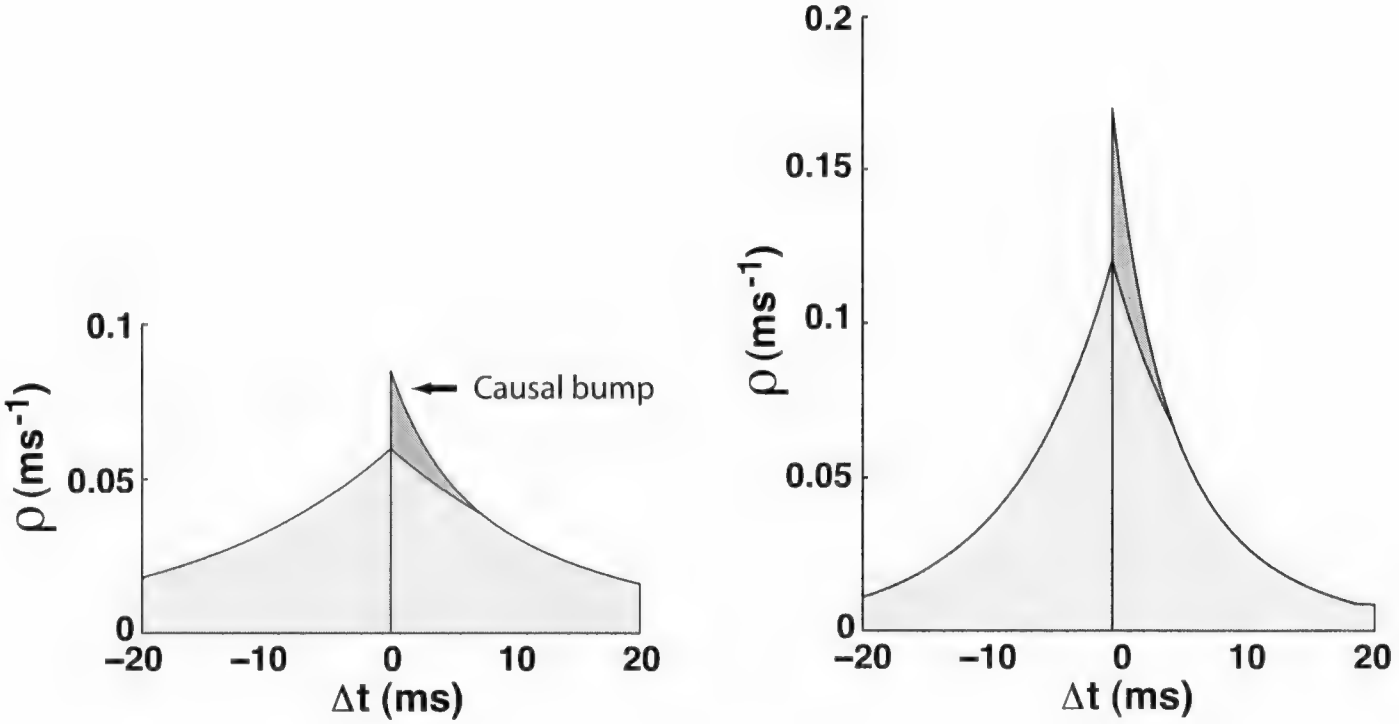
0.1
0.05
0
ρ (ms-1)
Causal bump
−20
−10
0
10
20
Δt (ms)
0.2
0.15
0.1
0.05
0
ρ (ms-1)
−20
−10
0
10
20
Δt (ms)

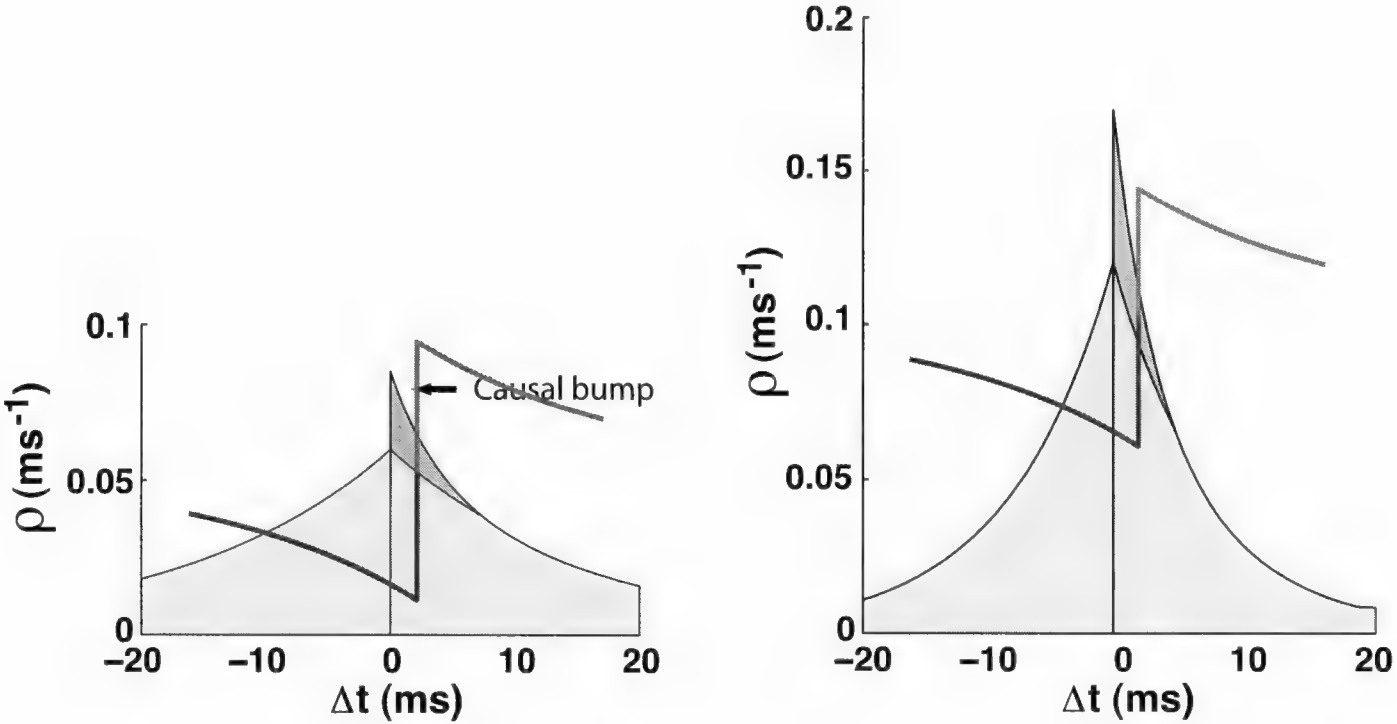
0.1
0.05
0
ρ (ms-1)
Causal bump
−20
−10
0
10
20
Δt (ms)
0.2
0.15
0.1
0.05
0
ρ (ms-1)
−20
−10
0
10
20
Δt (ms)

a

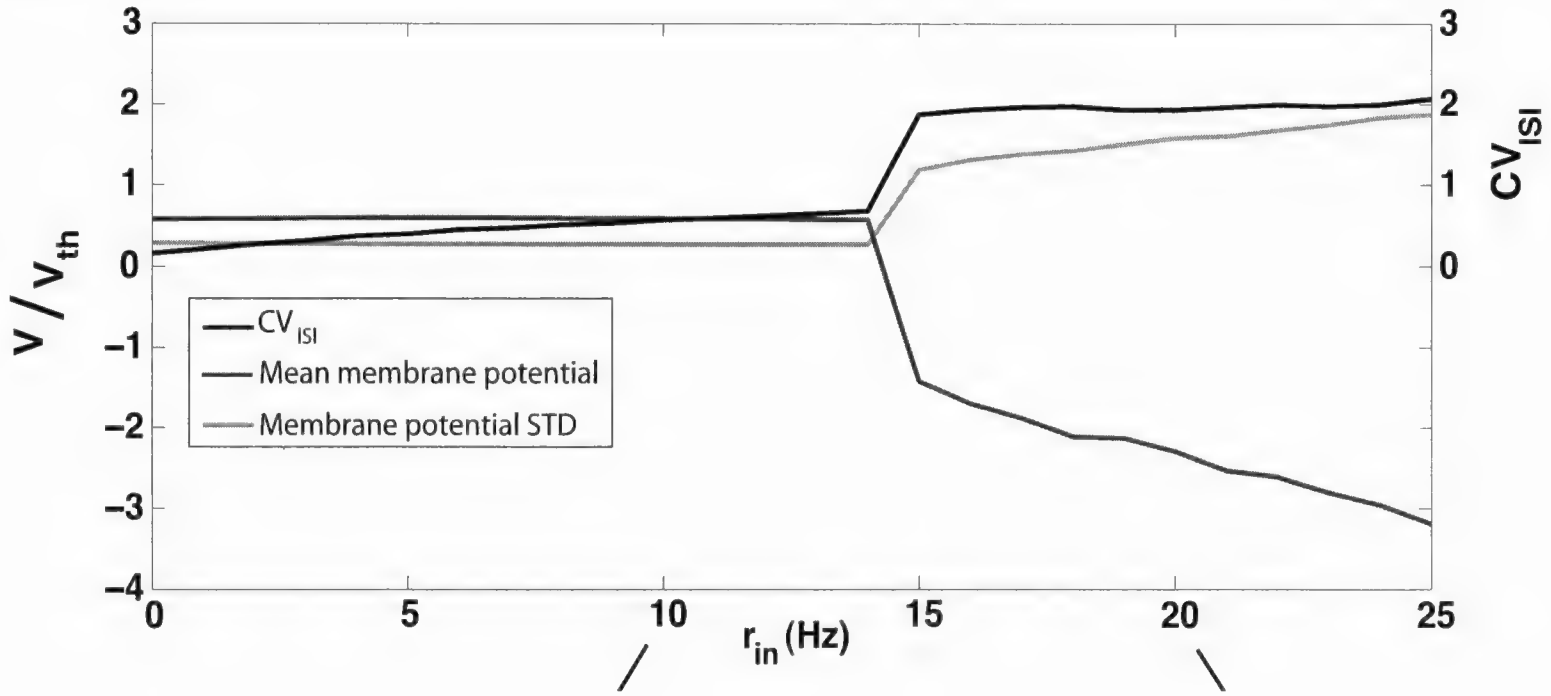
3
2
1
0
−1
−2
−3
−4
v / v_th
CV_ISI
Mean membrane potential
Membrane potential STD
0
5
10
15
20
25
r_in (Hz)

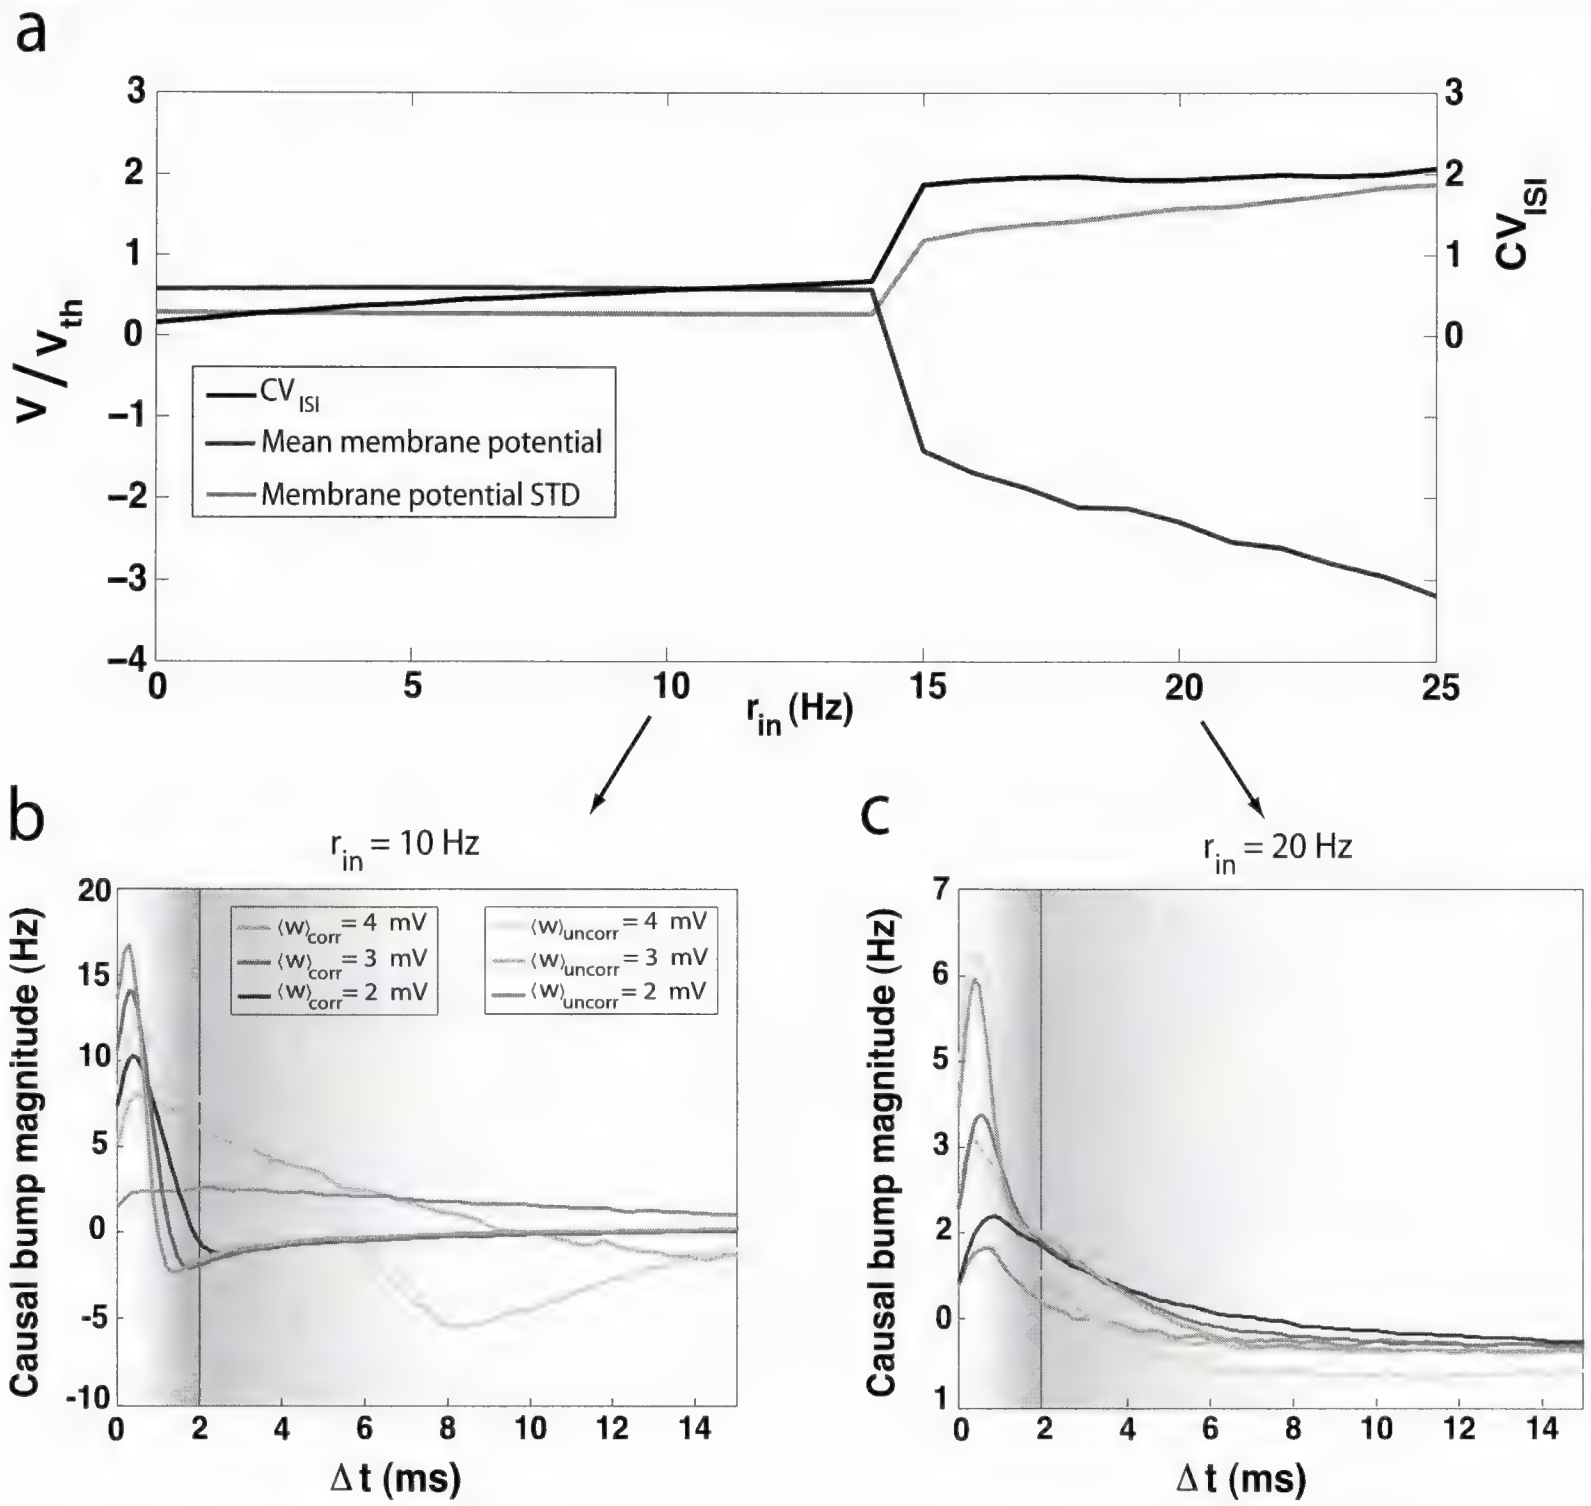
a
v / v_th
CV_ISI
Mean membrane potential
Membrane potential STD
r_in (Hz)
b
r_in = 10 Hz
Causal bump magnitude (Hz)
⟨w⟩_corr = 4 mV
⟨w⟩_corr = 3 mV
⟨w⟩_corr = 2 mV
⟨w⟩_uncorr = 4 mV
⟨w⟩_uncorr = 3 mV
⟨w⟩_uncorr = 2 mV
Δt (ms)
c
r_in = 20 Hz
Causal bump magnitude (Hz)
Δt (ms)

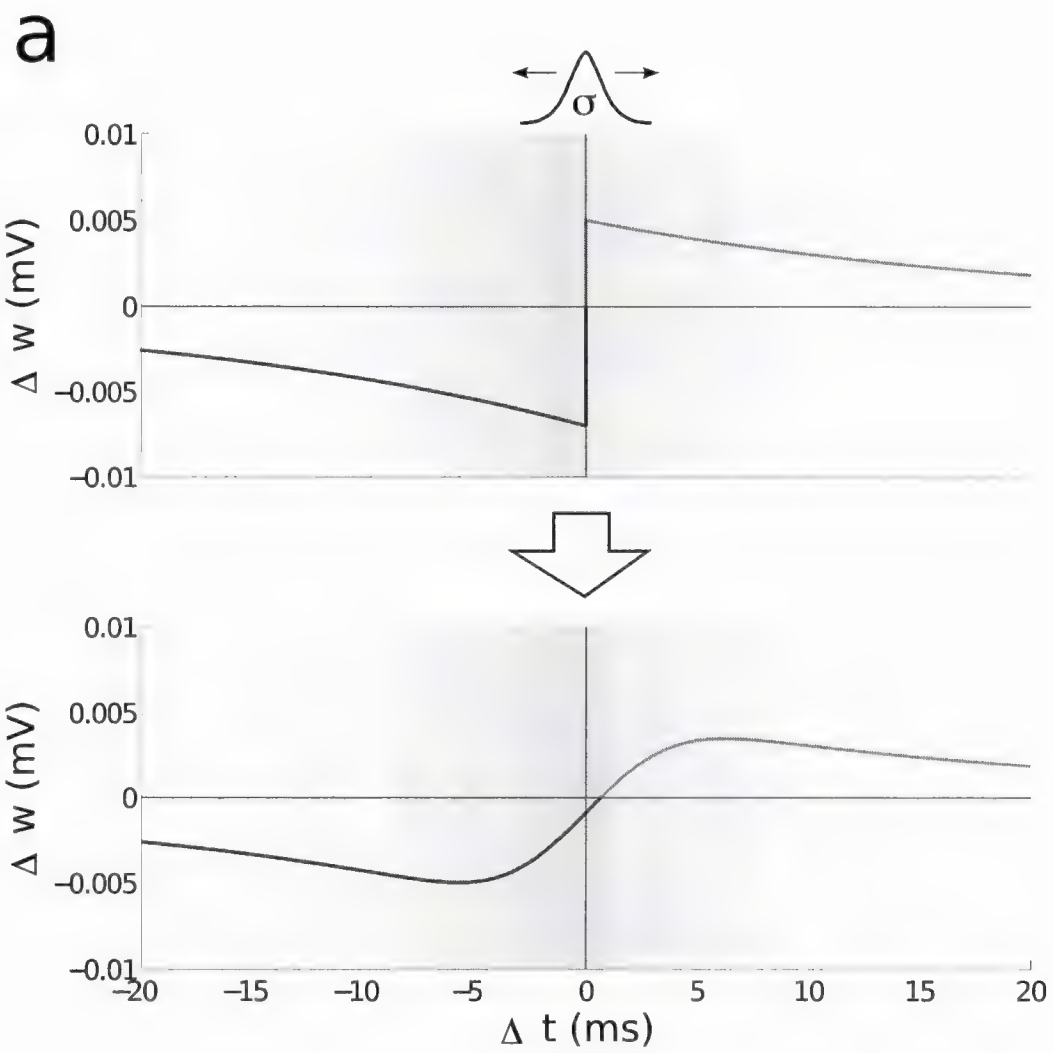
a
σ
0.01
0.005
0
−0.005
−0.01
Δ w (mV)
0.01
0.005
0
−0.005
−0.01
Δ w (mV)
−20
−15
−10
−5
0
5
10
15
20
Δ t (ms)

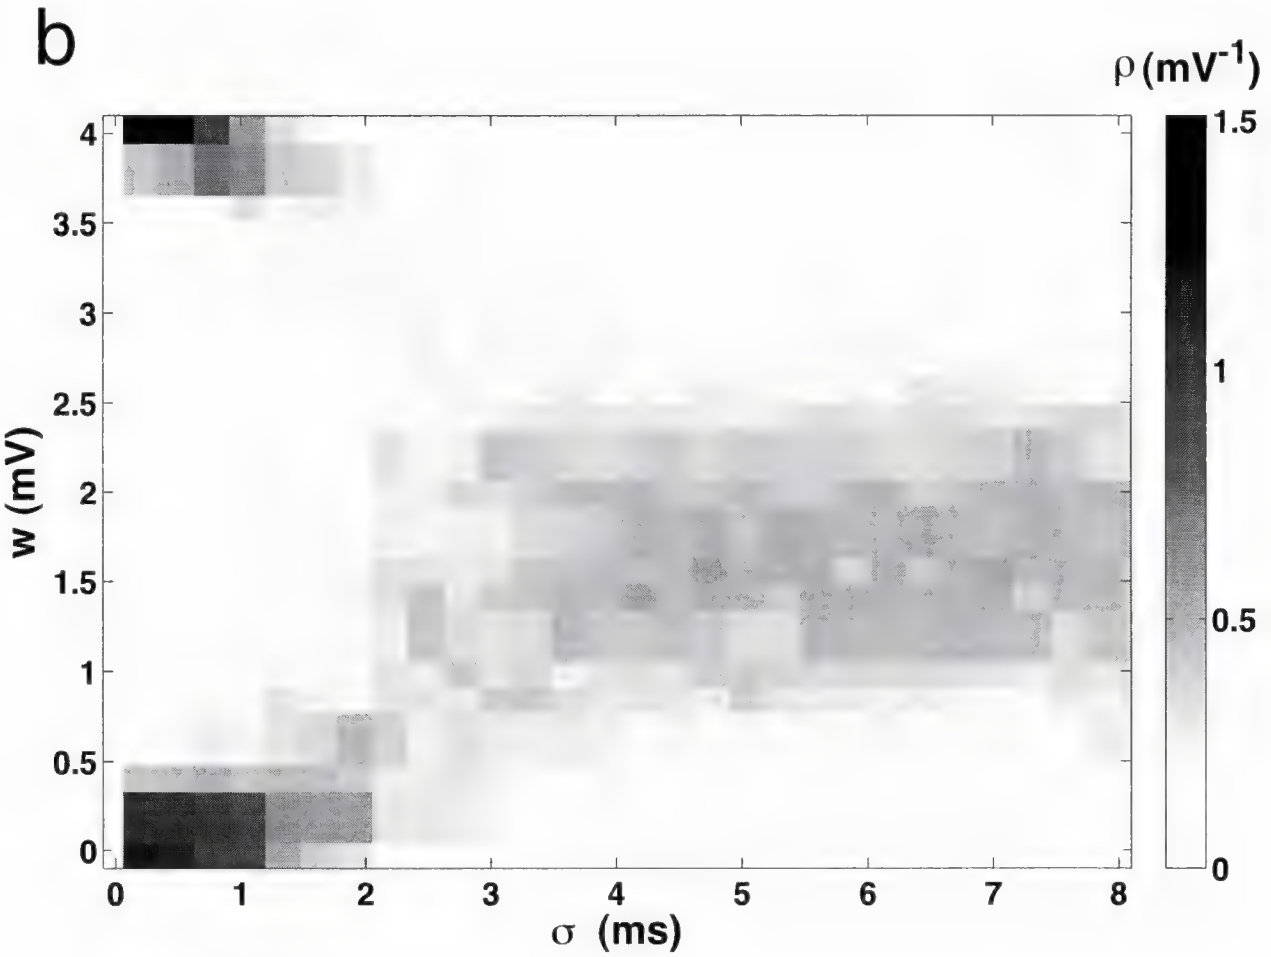
b
ρ (mV^{-1})
1.5
1
0.5
0
w (mV)
4
3.5
3
2.5
2
1.5
1
0.5
0
0
1
2
3
4
5
6
7
8
σ (ms)

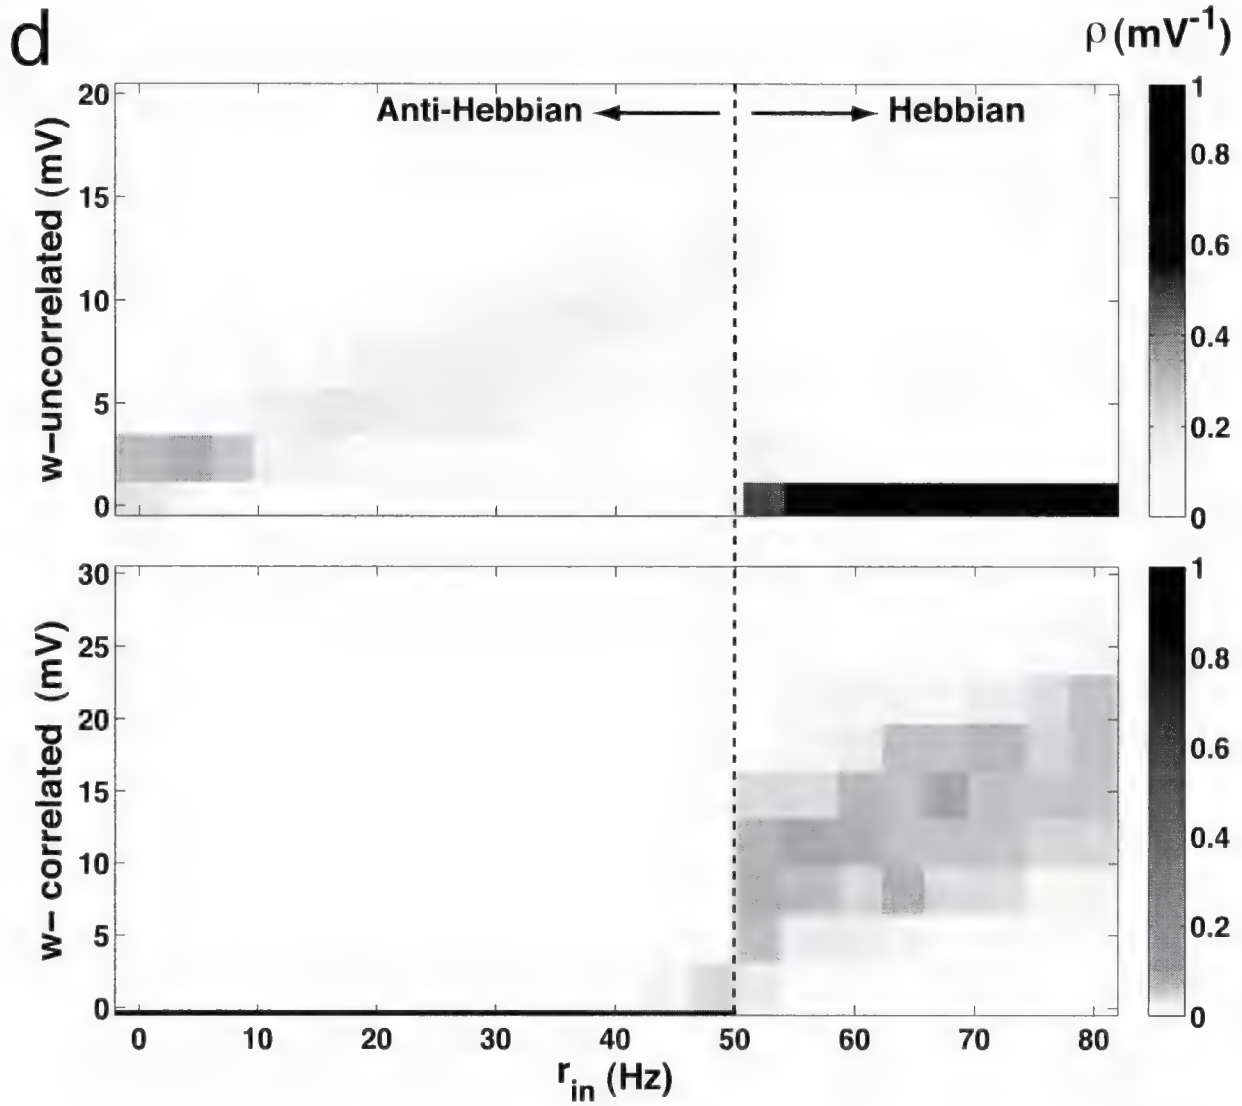
d
ρ (mV-1)
Anti-Hebbian
Hebbian
w–uncorrelated (mV)
w– correlated (mV)
rin (Hz)
0
5
10
15
20
25
30
40
50
60
70
80
0.2
0.4
0.6
0.8
1

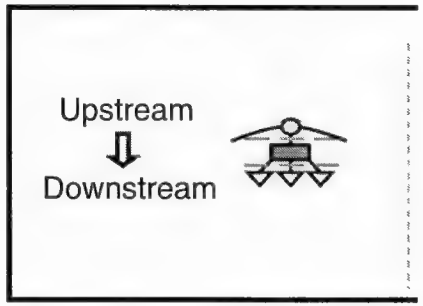

Evolutionary expansion and anatomical specialization of synapse proteome complexity
Nature Neuroscience 2008

Richard D Emes, Andrew J Pocklington, Christopher N G Anderson, Alex Bayes, Mark O Collins, Catherine A Vickers, Mike D R Croning, Bilal R Malik, Jyoti S Choudhary, J Douglas Armstrong & Seth G N Grant

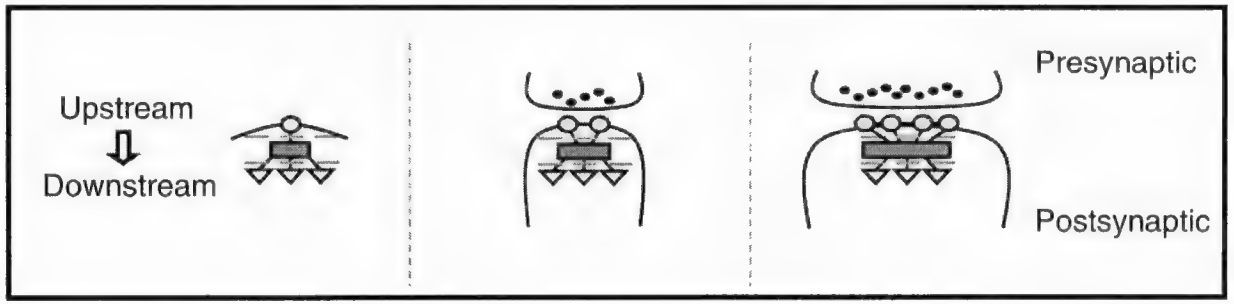

Evolutionary expansion and anatomical specialization of synapse proteome complexity
Nature Neuroscience 2008

Richard D Emes, Andrew J Pocklington, Christopher N G Anderson, Alex Bayes, Mark O Collins, Catherine A Vickers, Mike D R Croning, Bilal R Malik, Jyoti S Choudhary, J Douglas Armstrong & Seth G N Grant

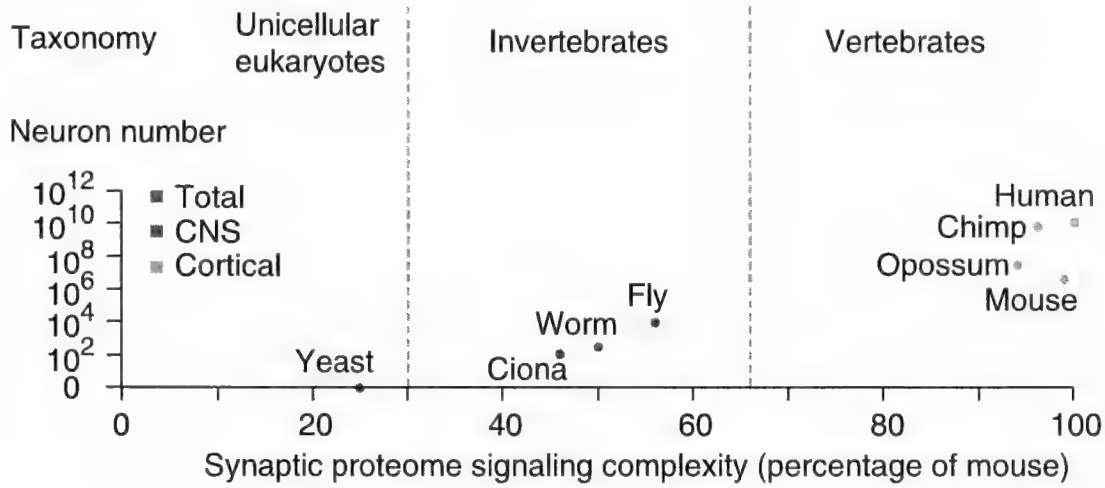

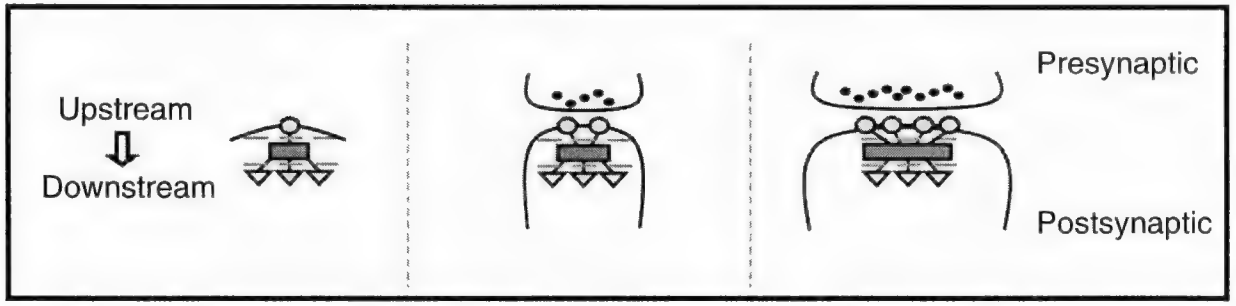

Evolutionary expansion and anatomical specialization of synapse proteome complexity
Nature Neuroscience 2008

Richard D Emes, Andrew J Pocklington, Christopher N G Anderson, Alex Bayes, Mark O Collins, Catherine A Vickers, Mike D R Croning, Bilal R Malik, Jyoti S Choudhary, J Douglas Armstrong & Seth G N Grant

Comprising over 1000 proteins, the macromolecular complexes of neurotransmitter receptors connected with the postsynaptic density (PSD) are perhaps the most complex molecular structures known in mammals.

Andrew J Pocklington, Mark Cumiskey, J Douglas Armstrong and Seth Grant
Molecular Systems Biology (2006)

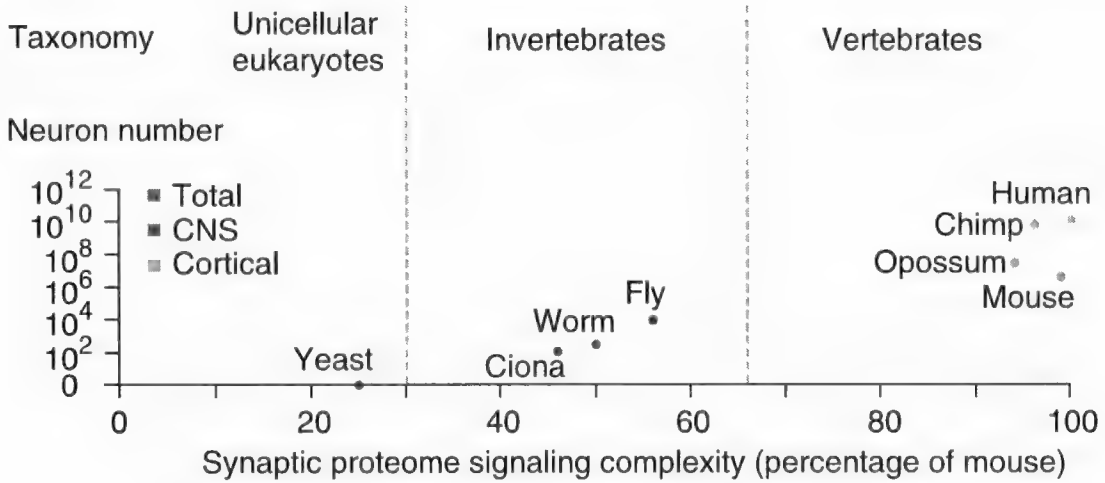

Evolutionary expansion and anatomical specialization of synapse proteome complexity
Nature Neuroscience 2008

Richard D Emes, Andrew J Pocklington, Christopher N G Anderson, Alex Bayes, Mark O Collins, Catherine A Vickers, Mike D R Croning, Bilal R Malik, Jyoti S Choudhary, J Douglas Armstrong & Seth G N Grant

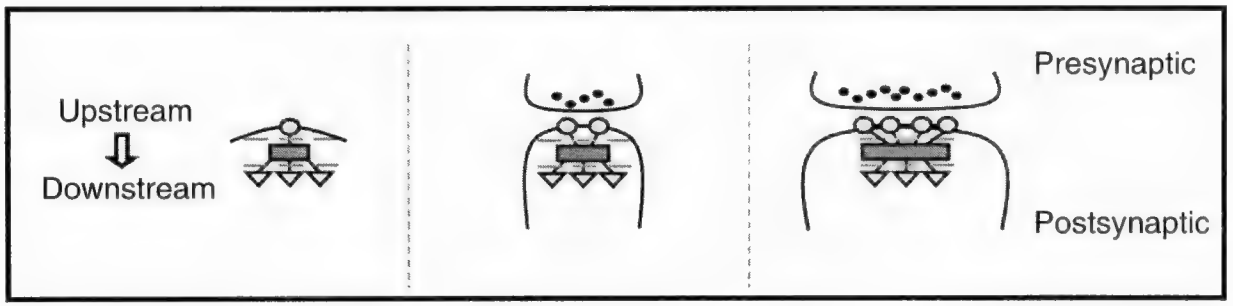

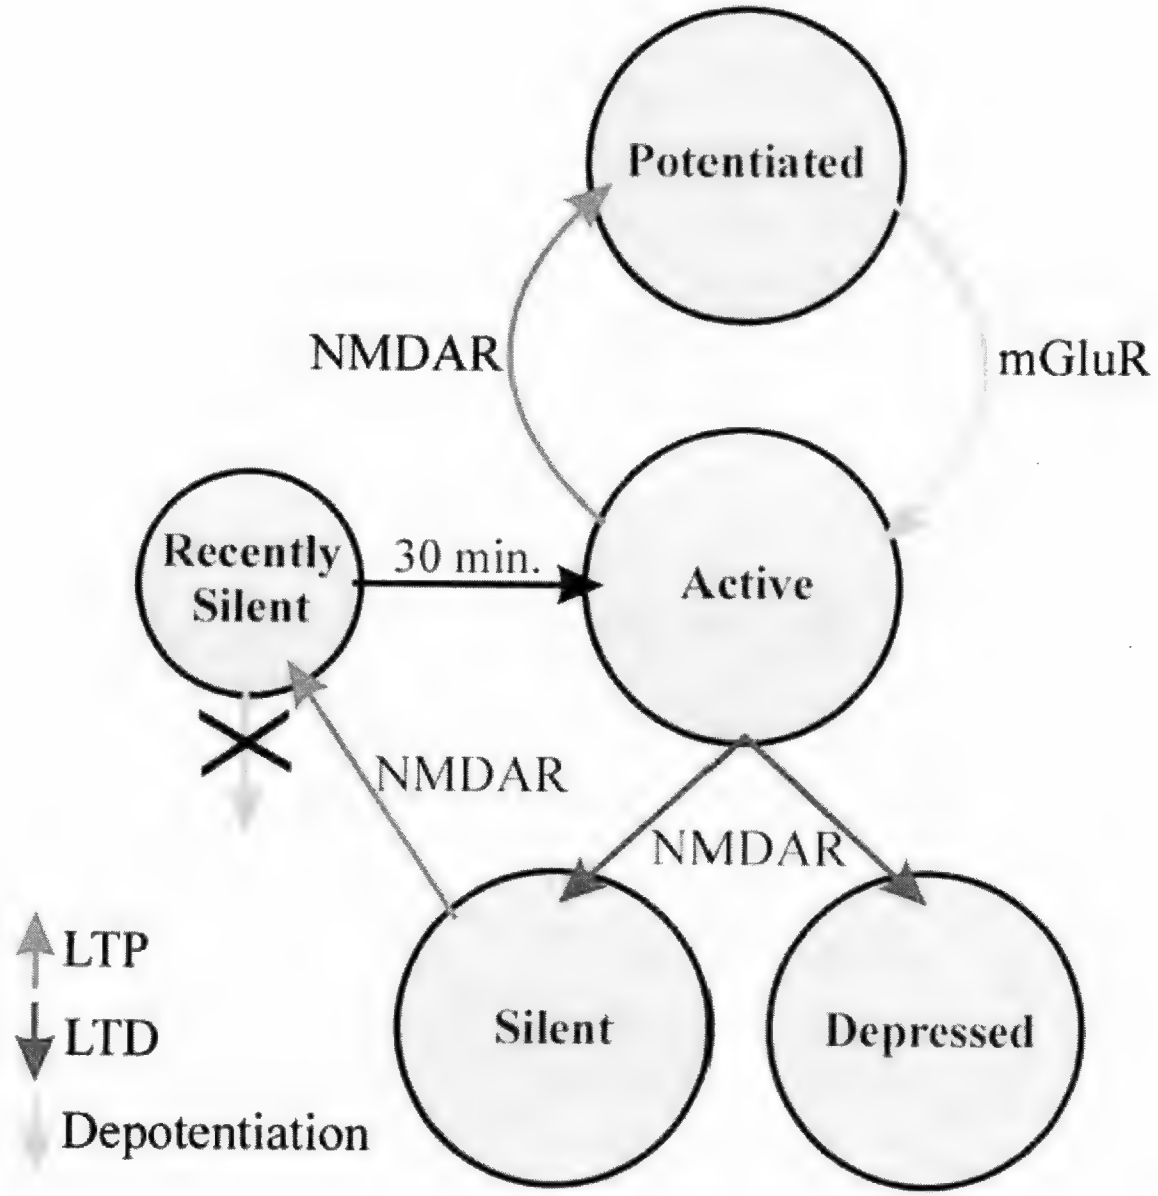

Montgomery & Madison, 2002

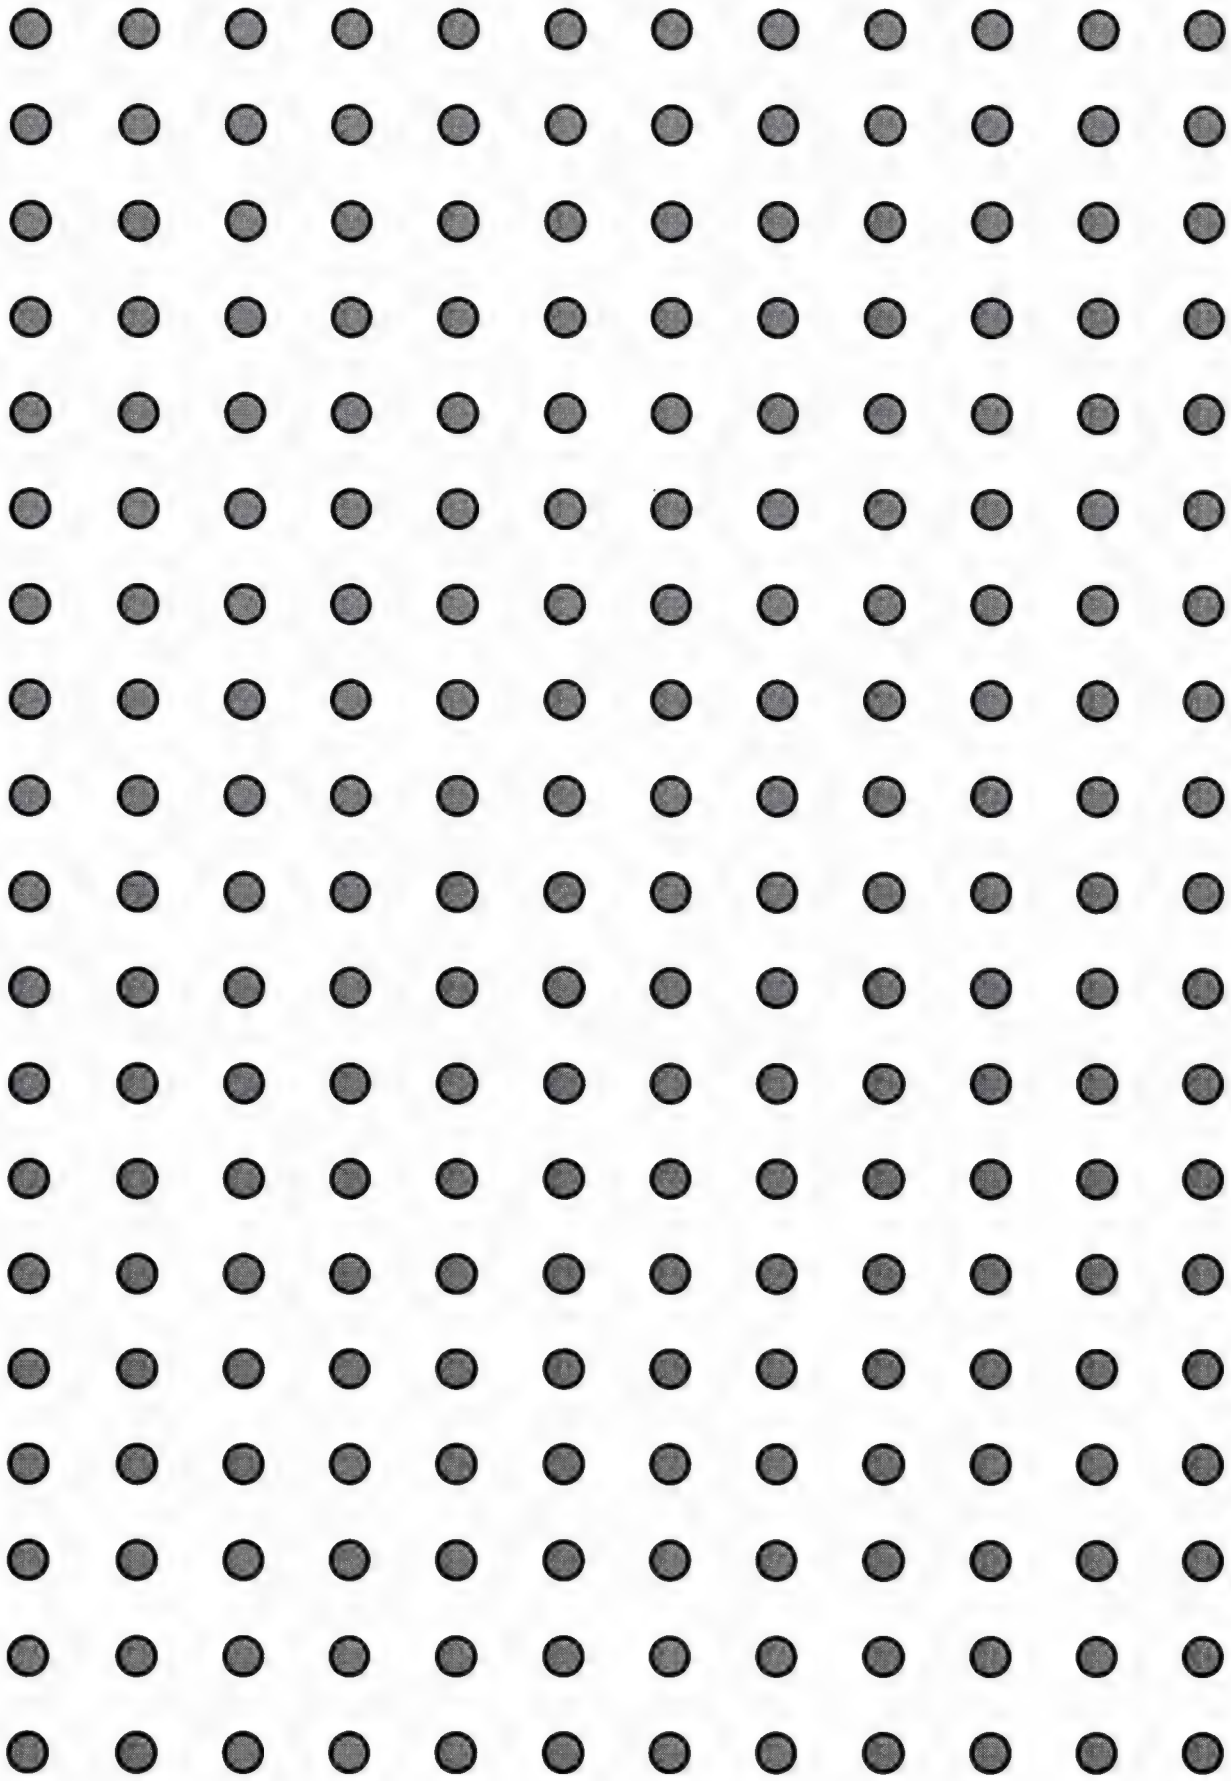

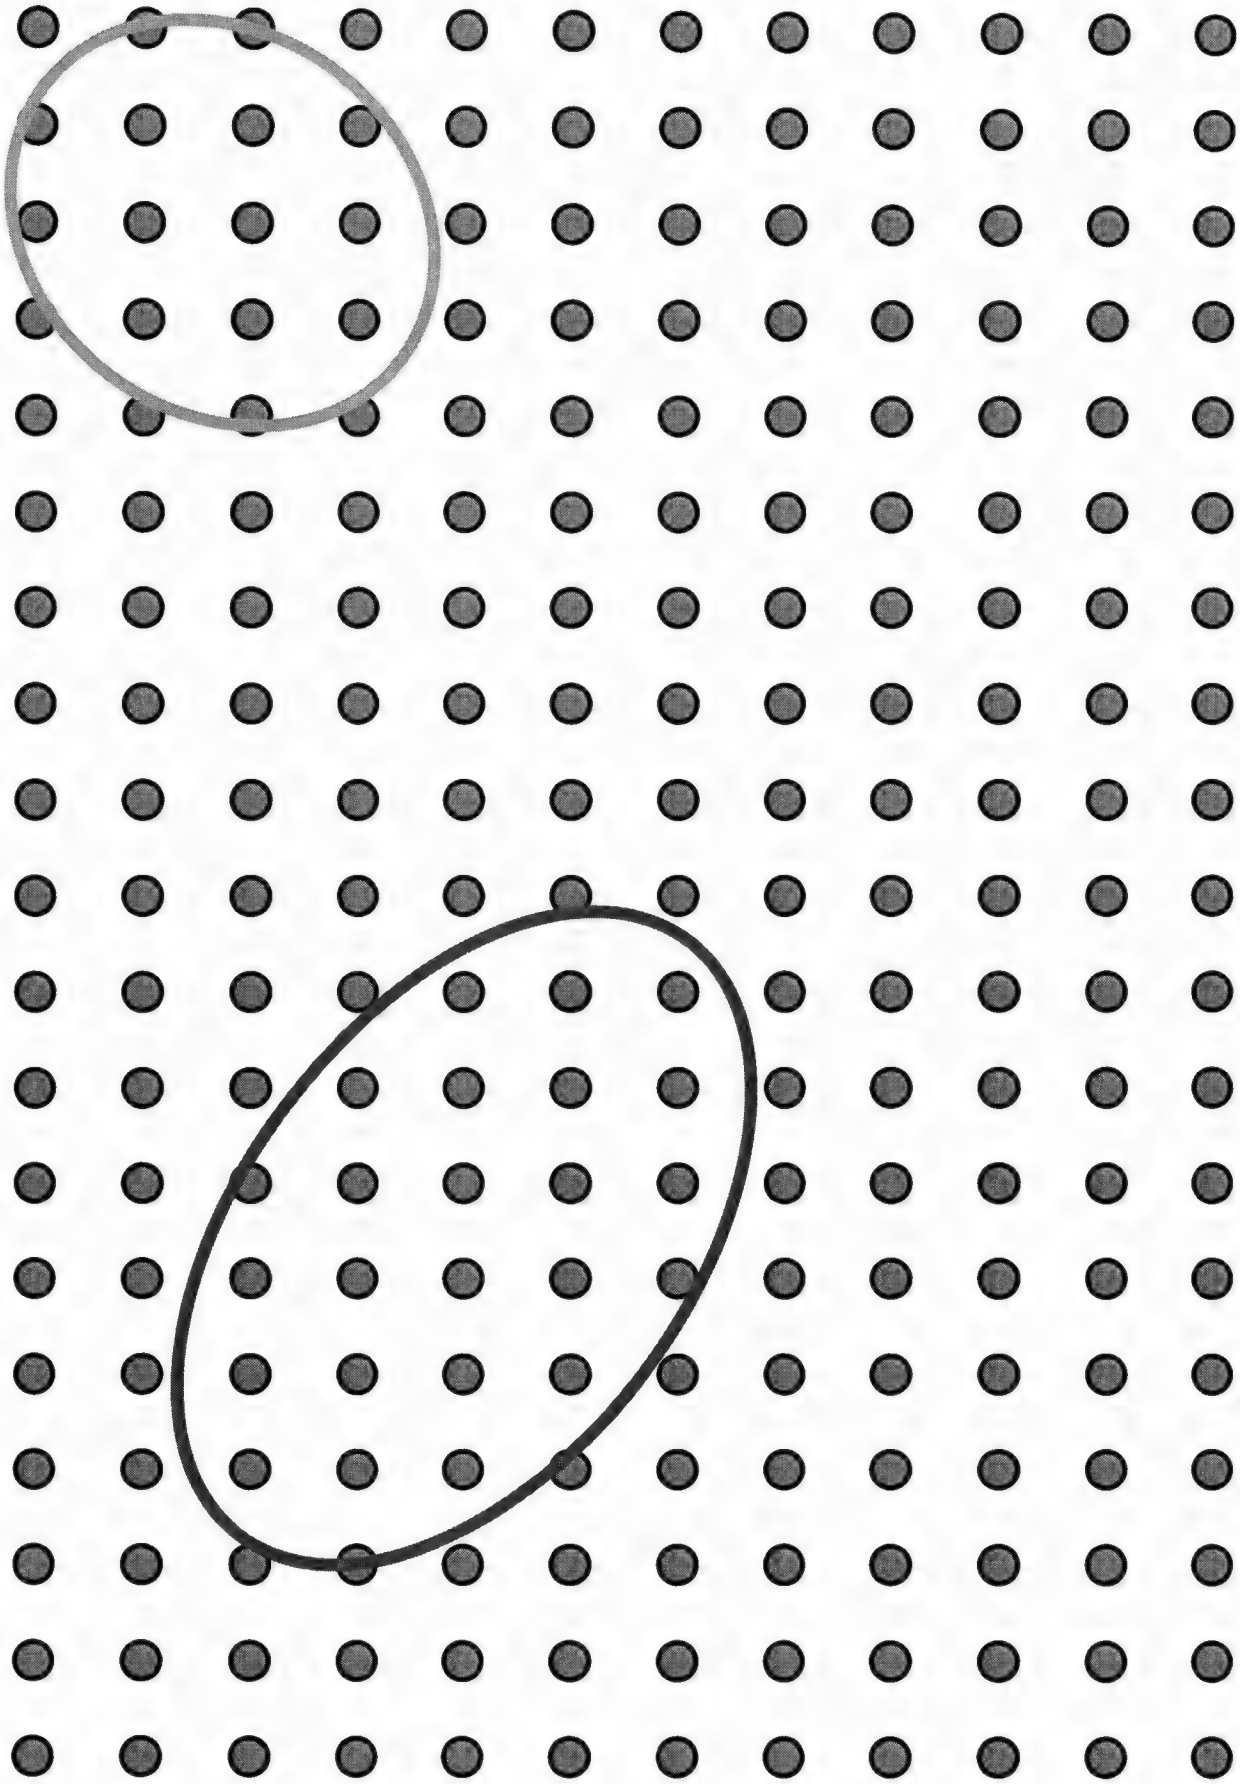

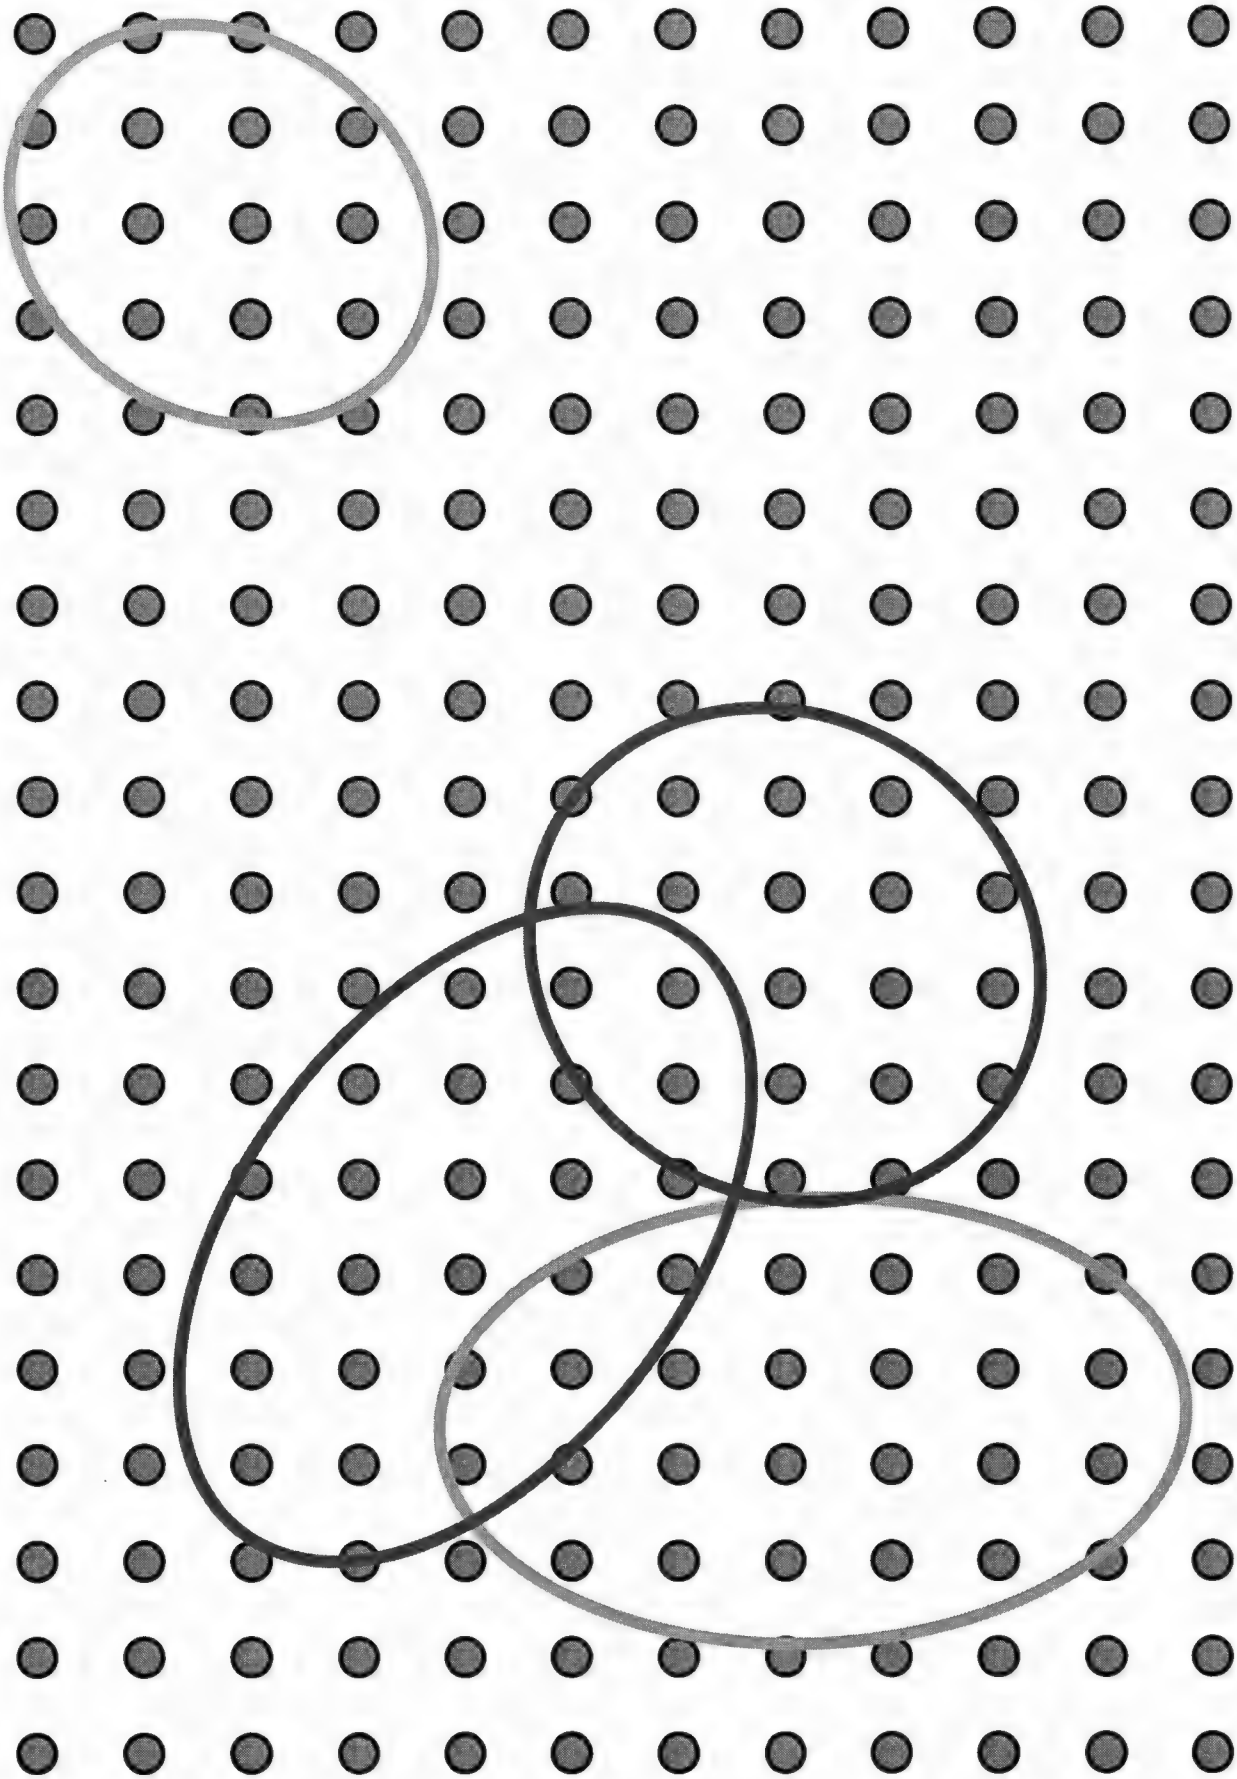

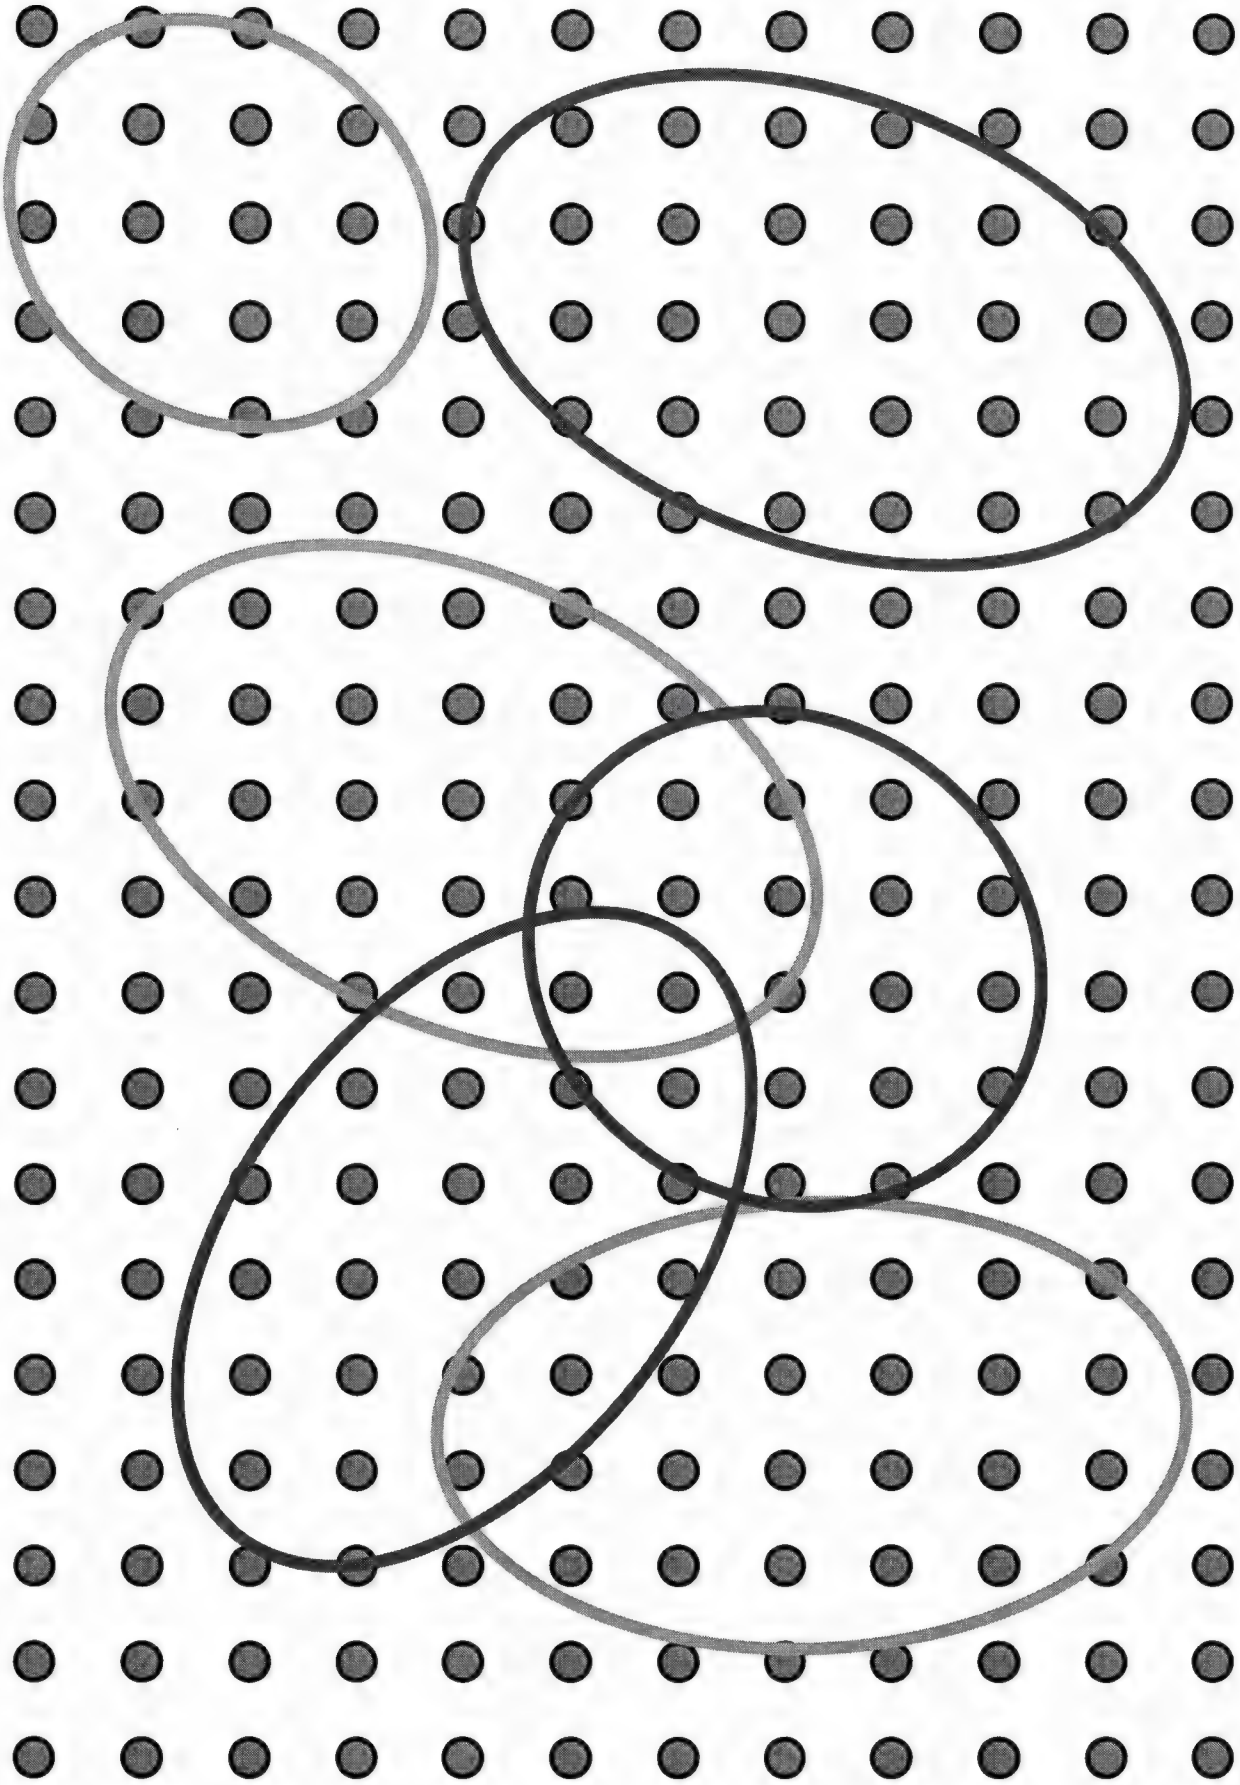

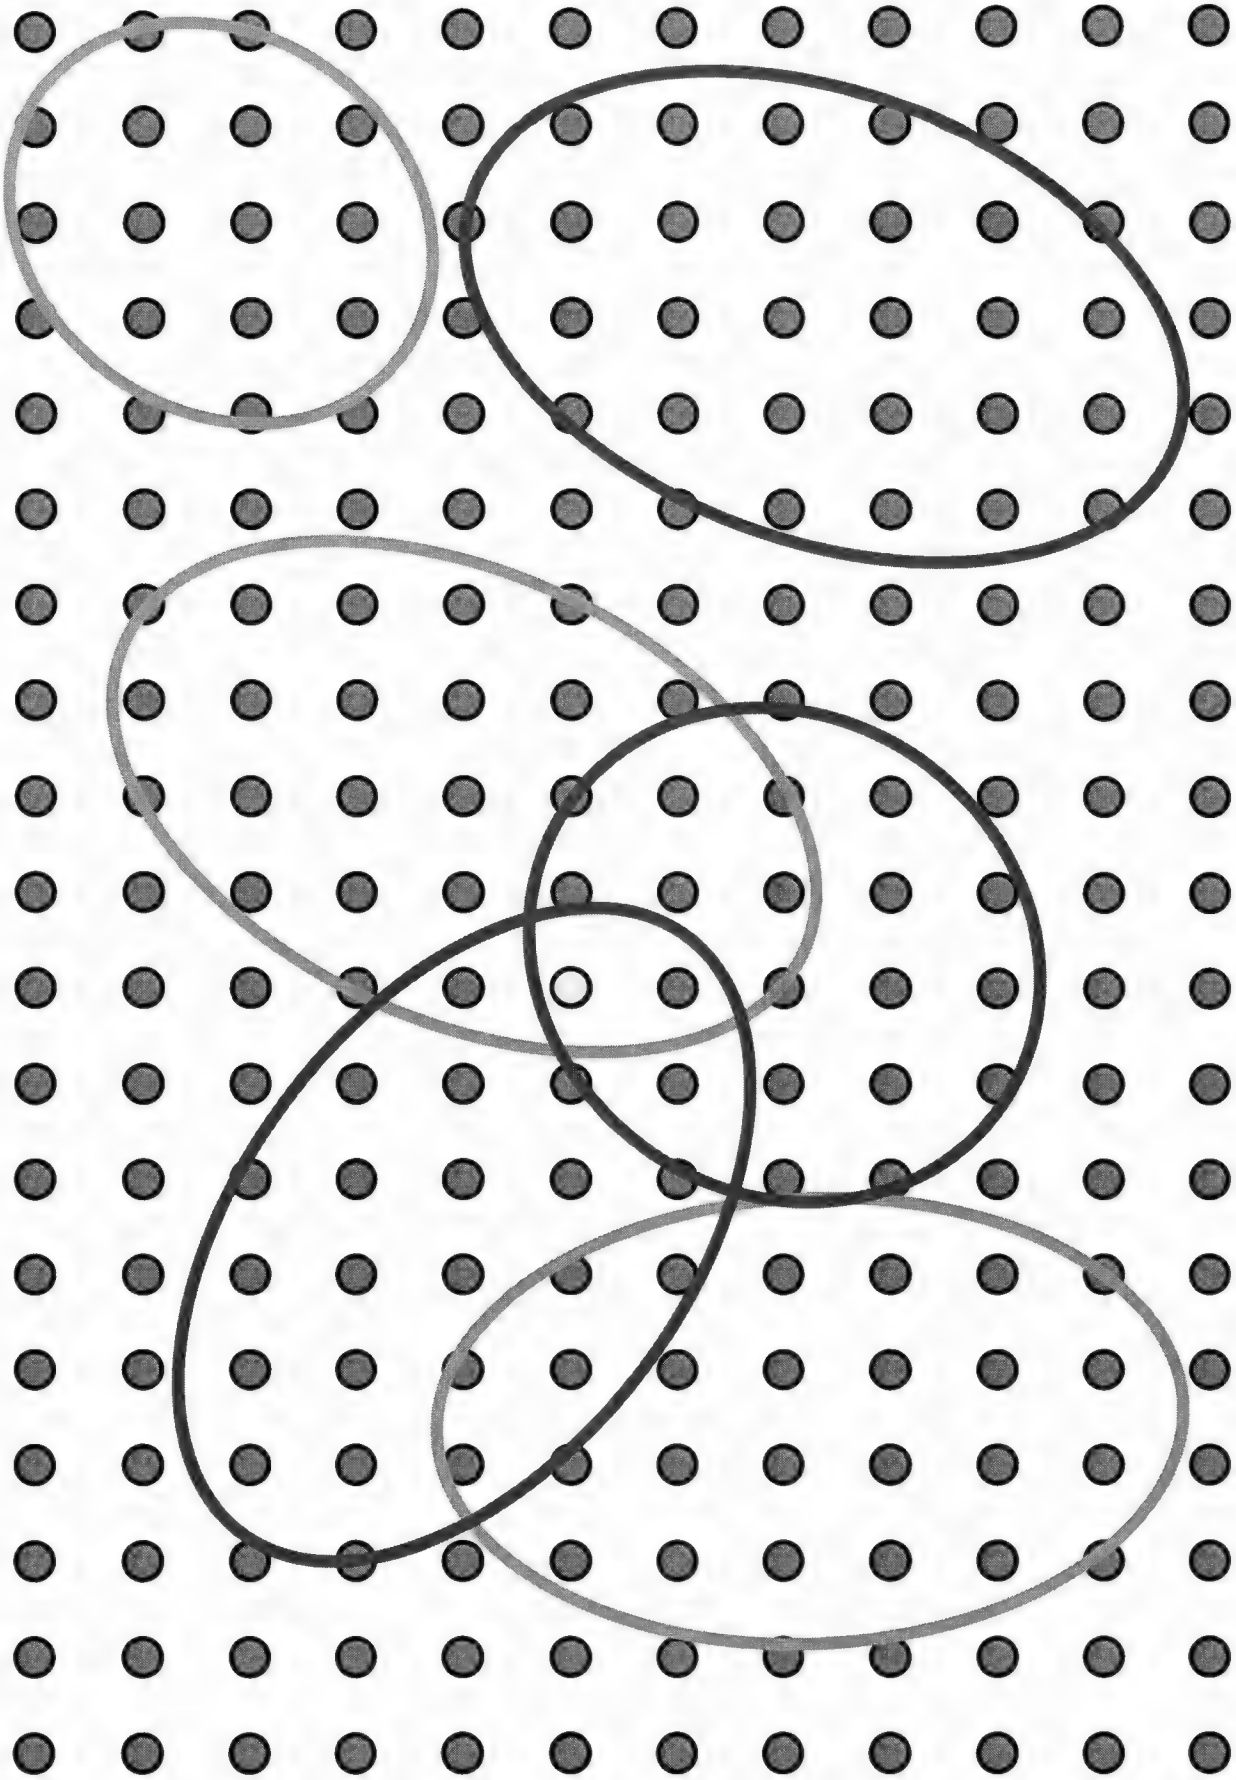

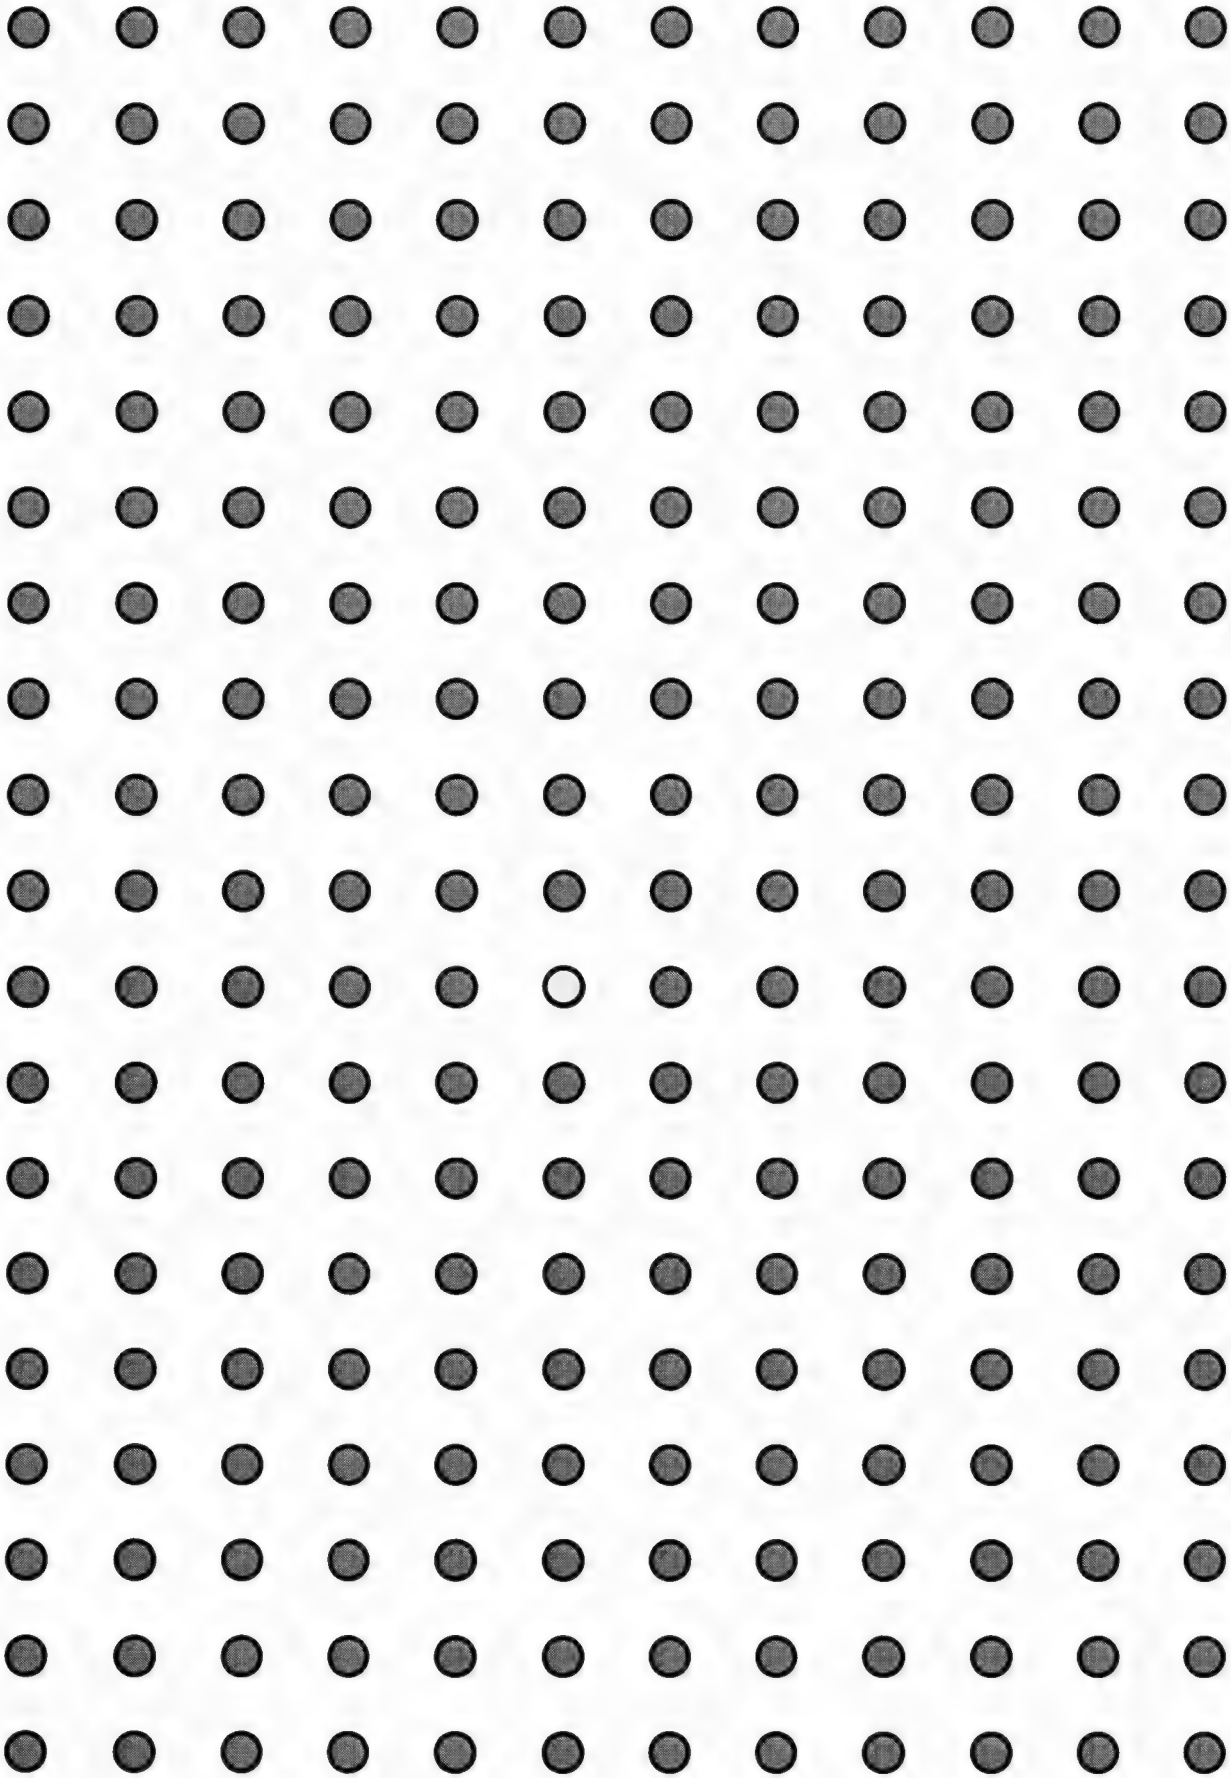

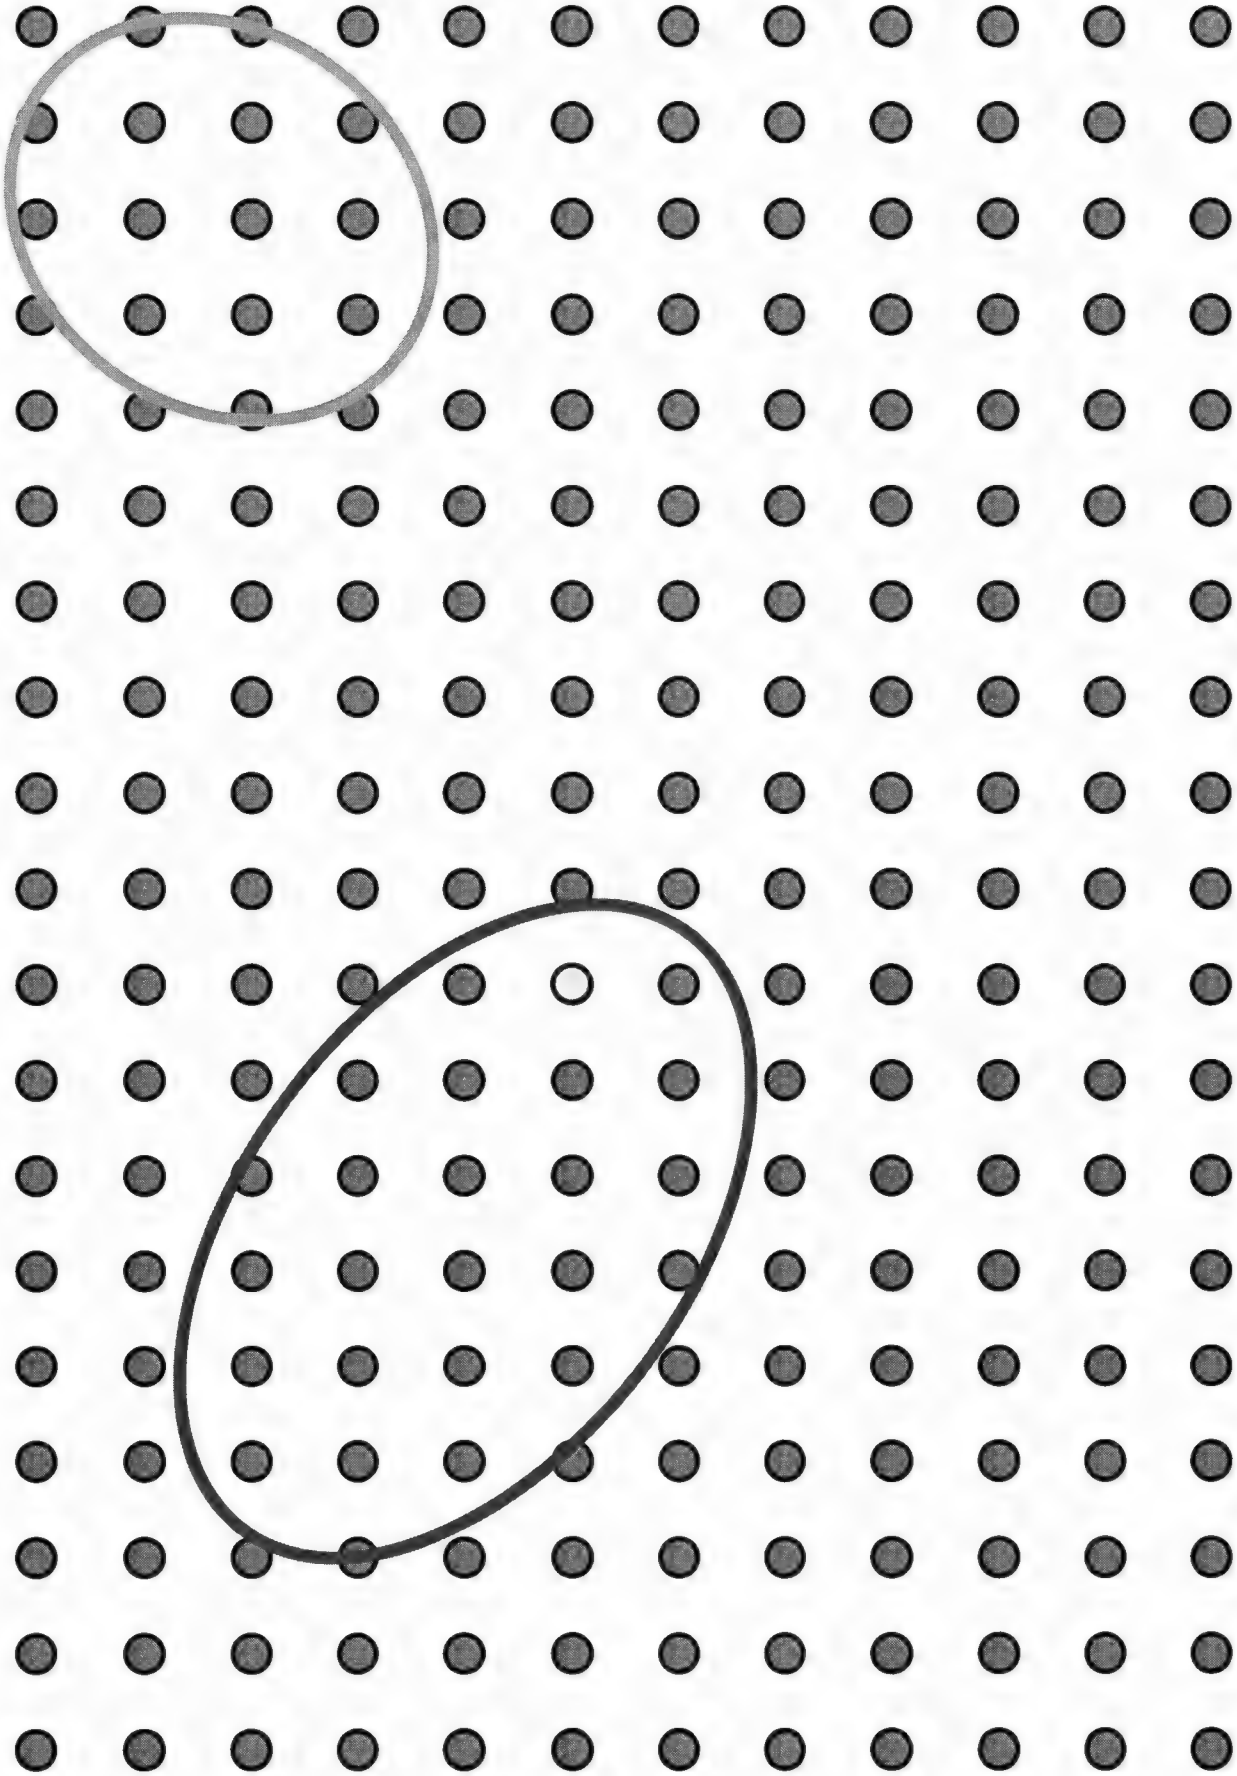

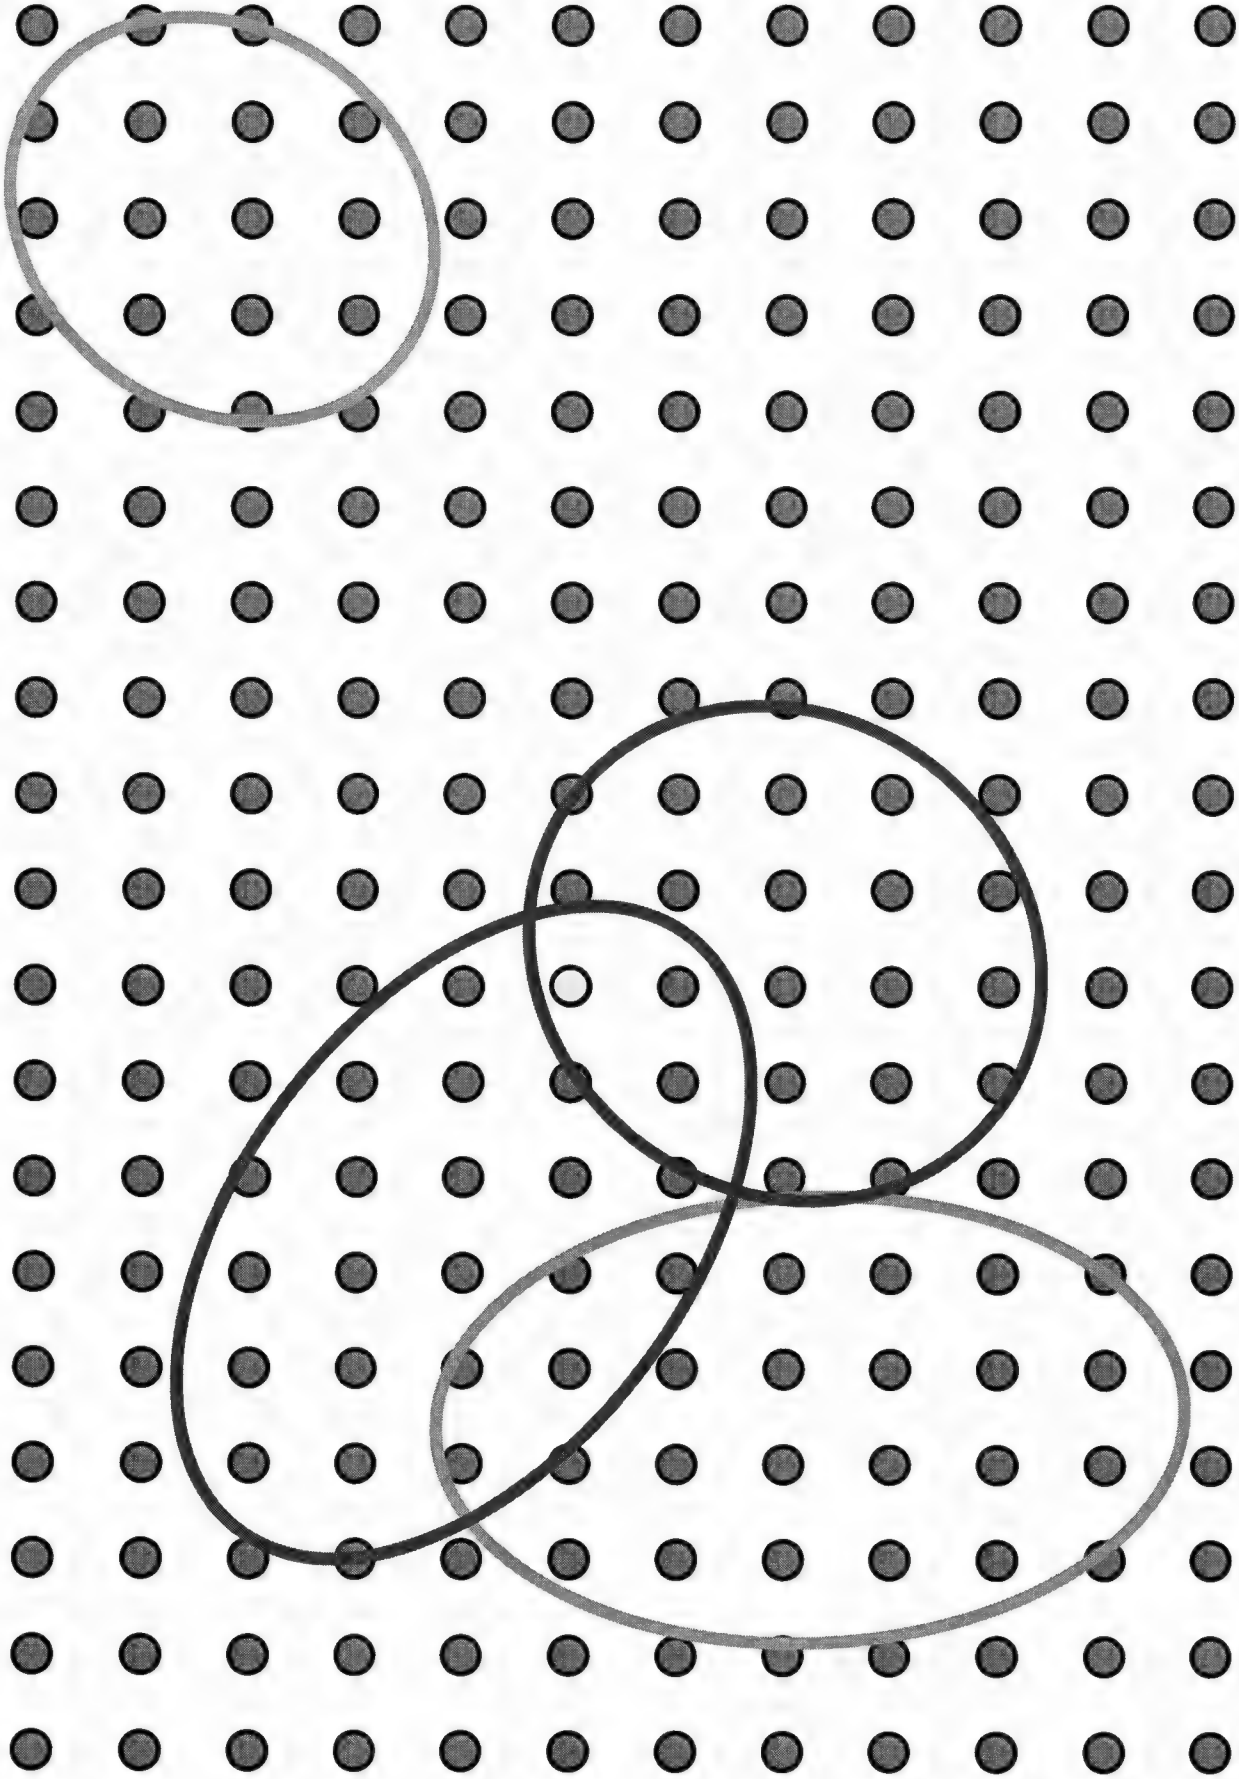

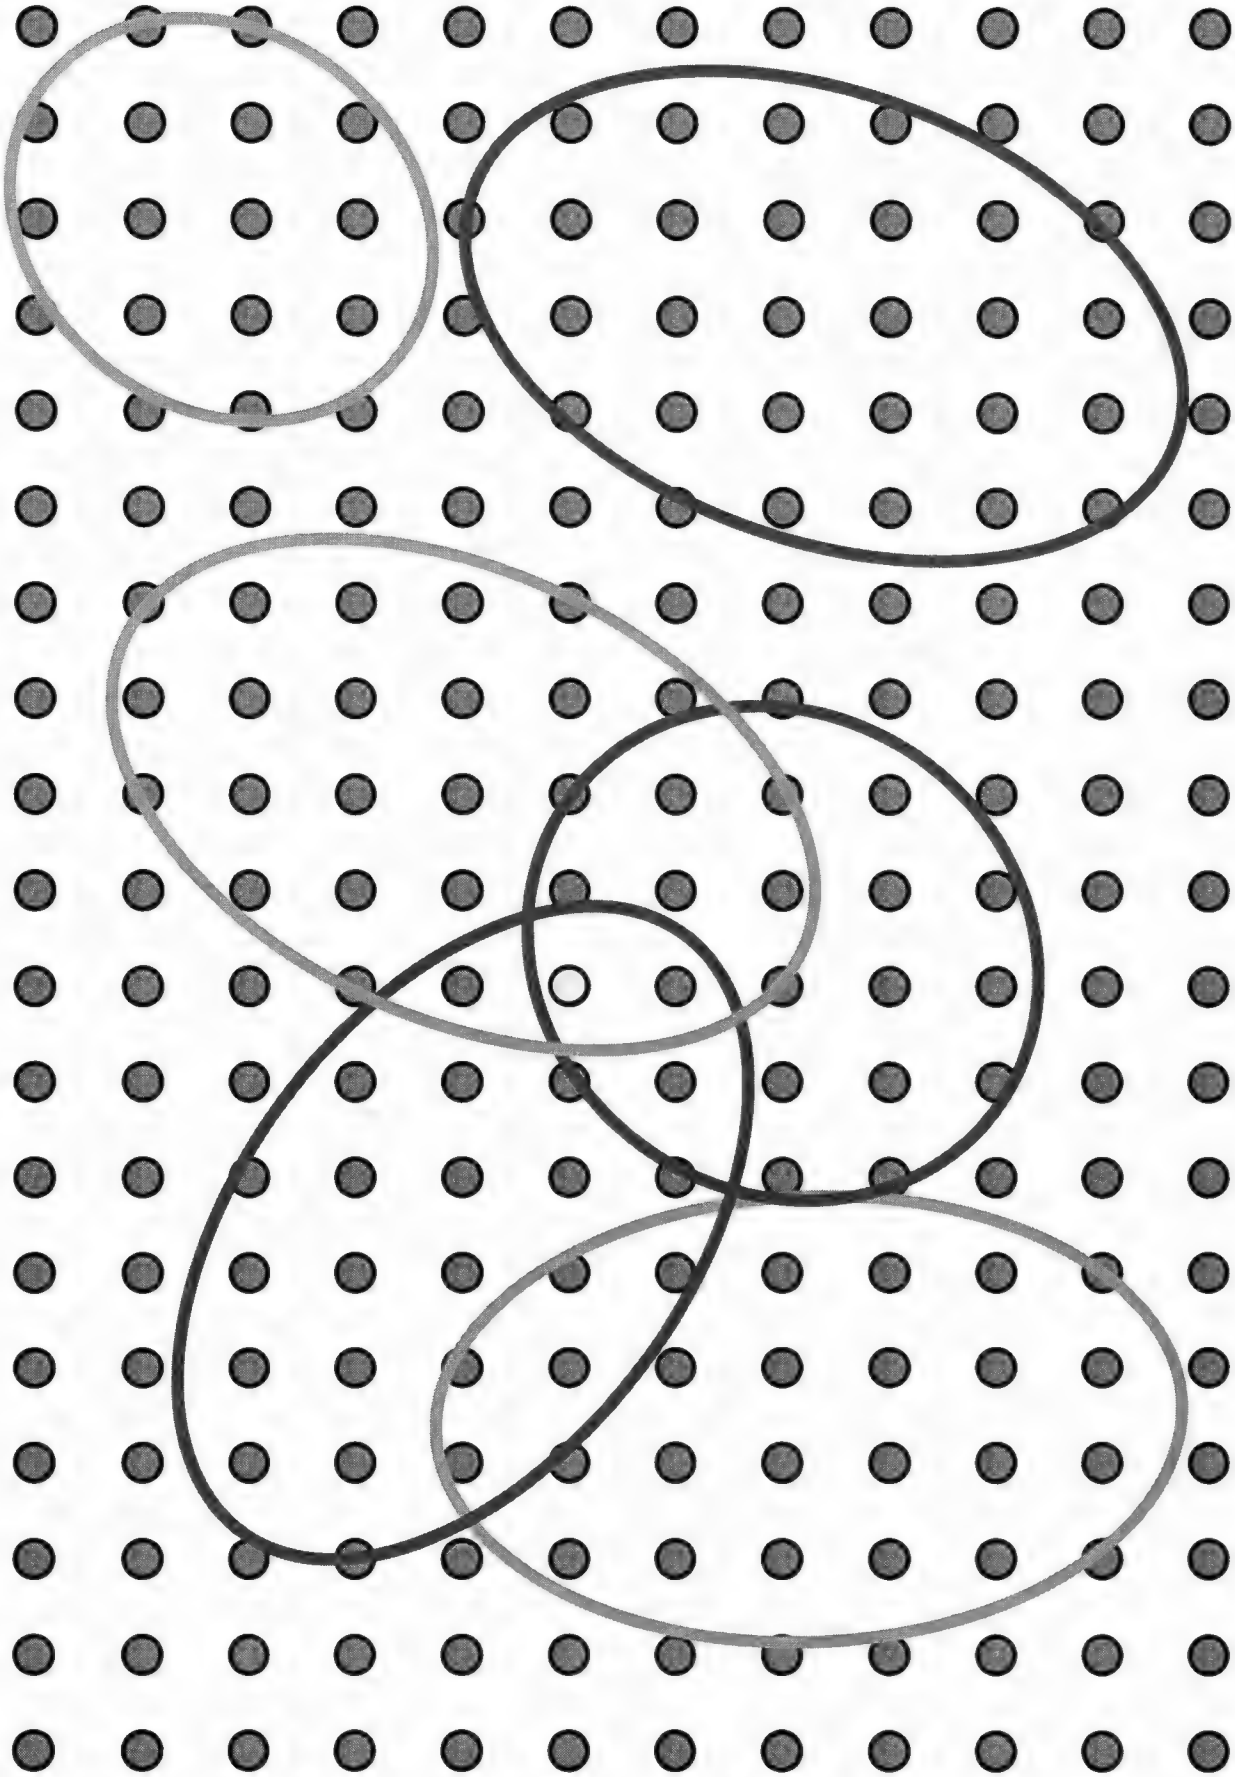

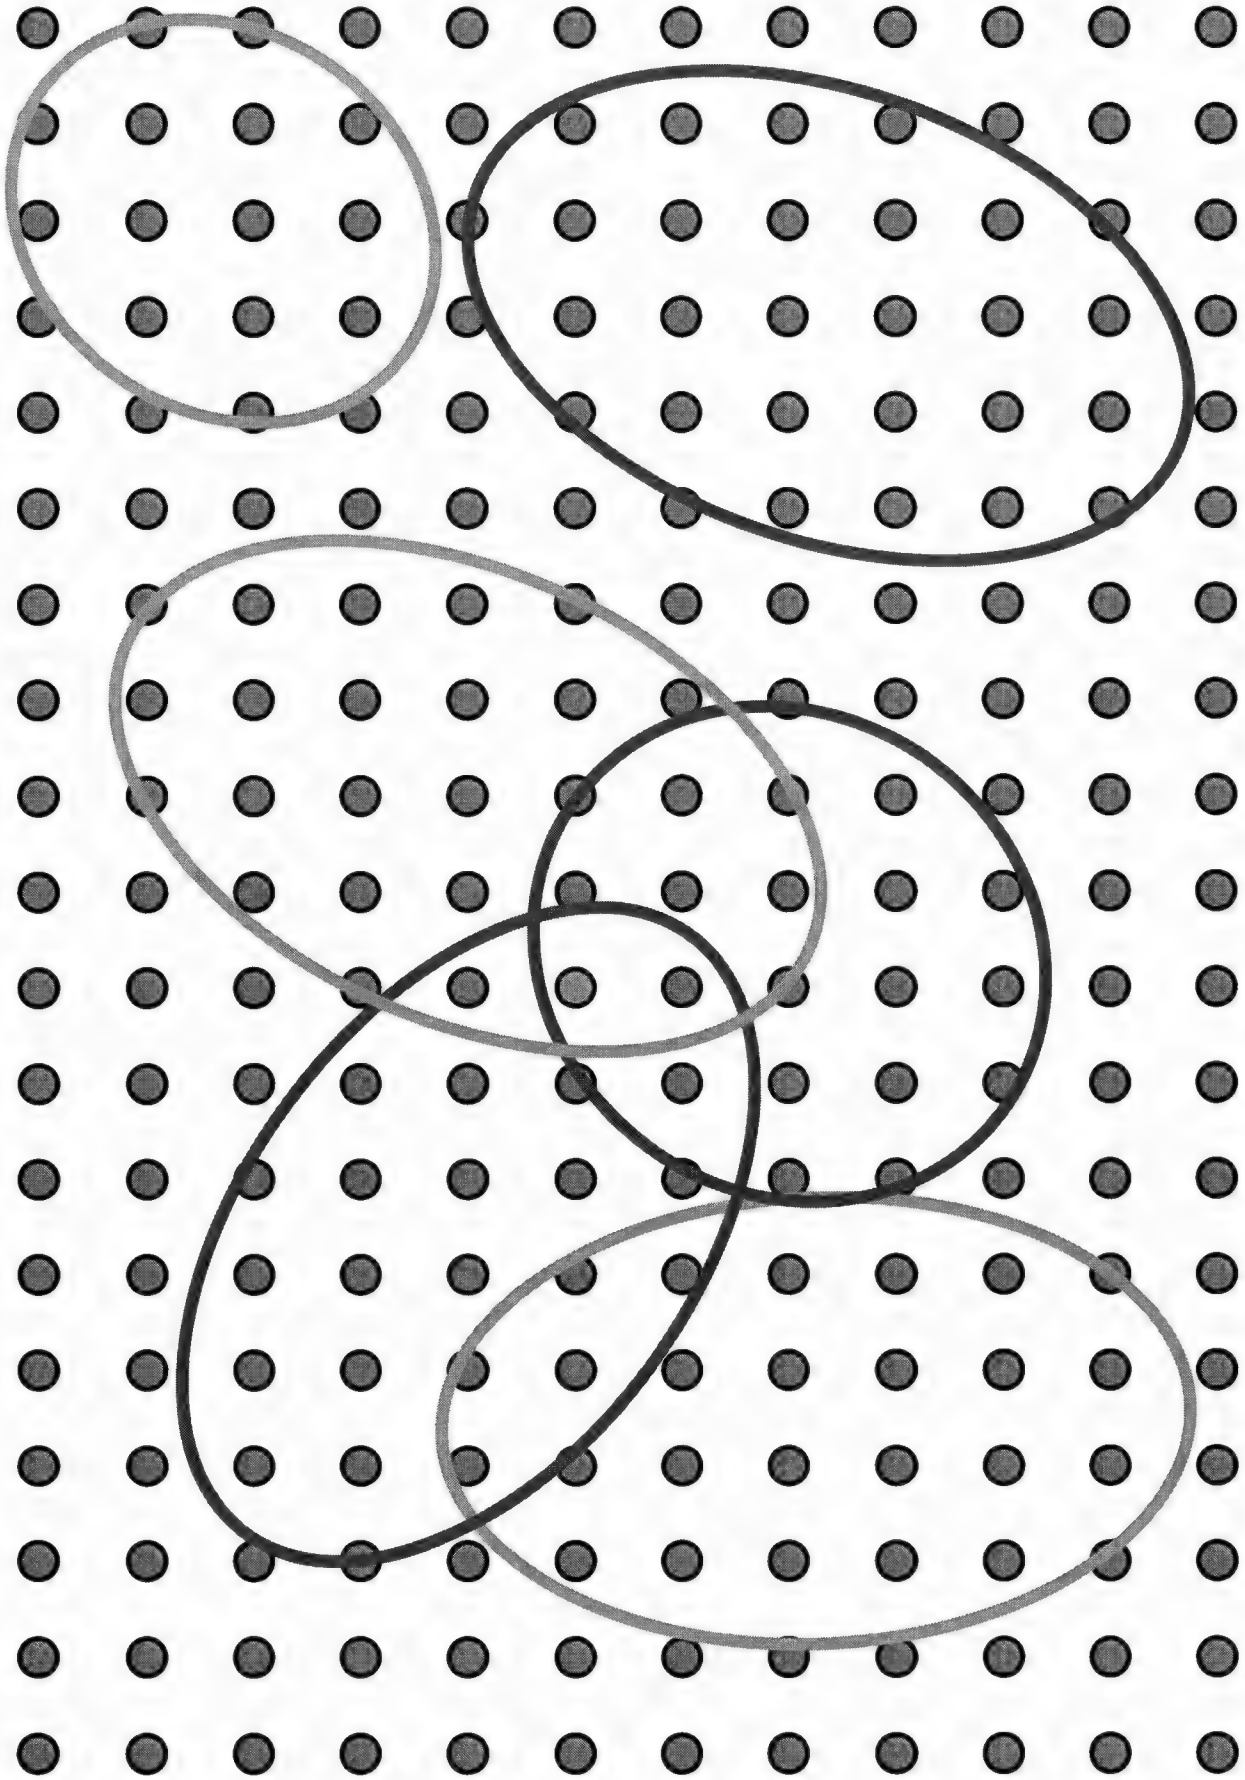

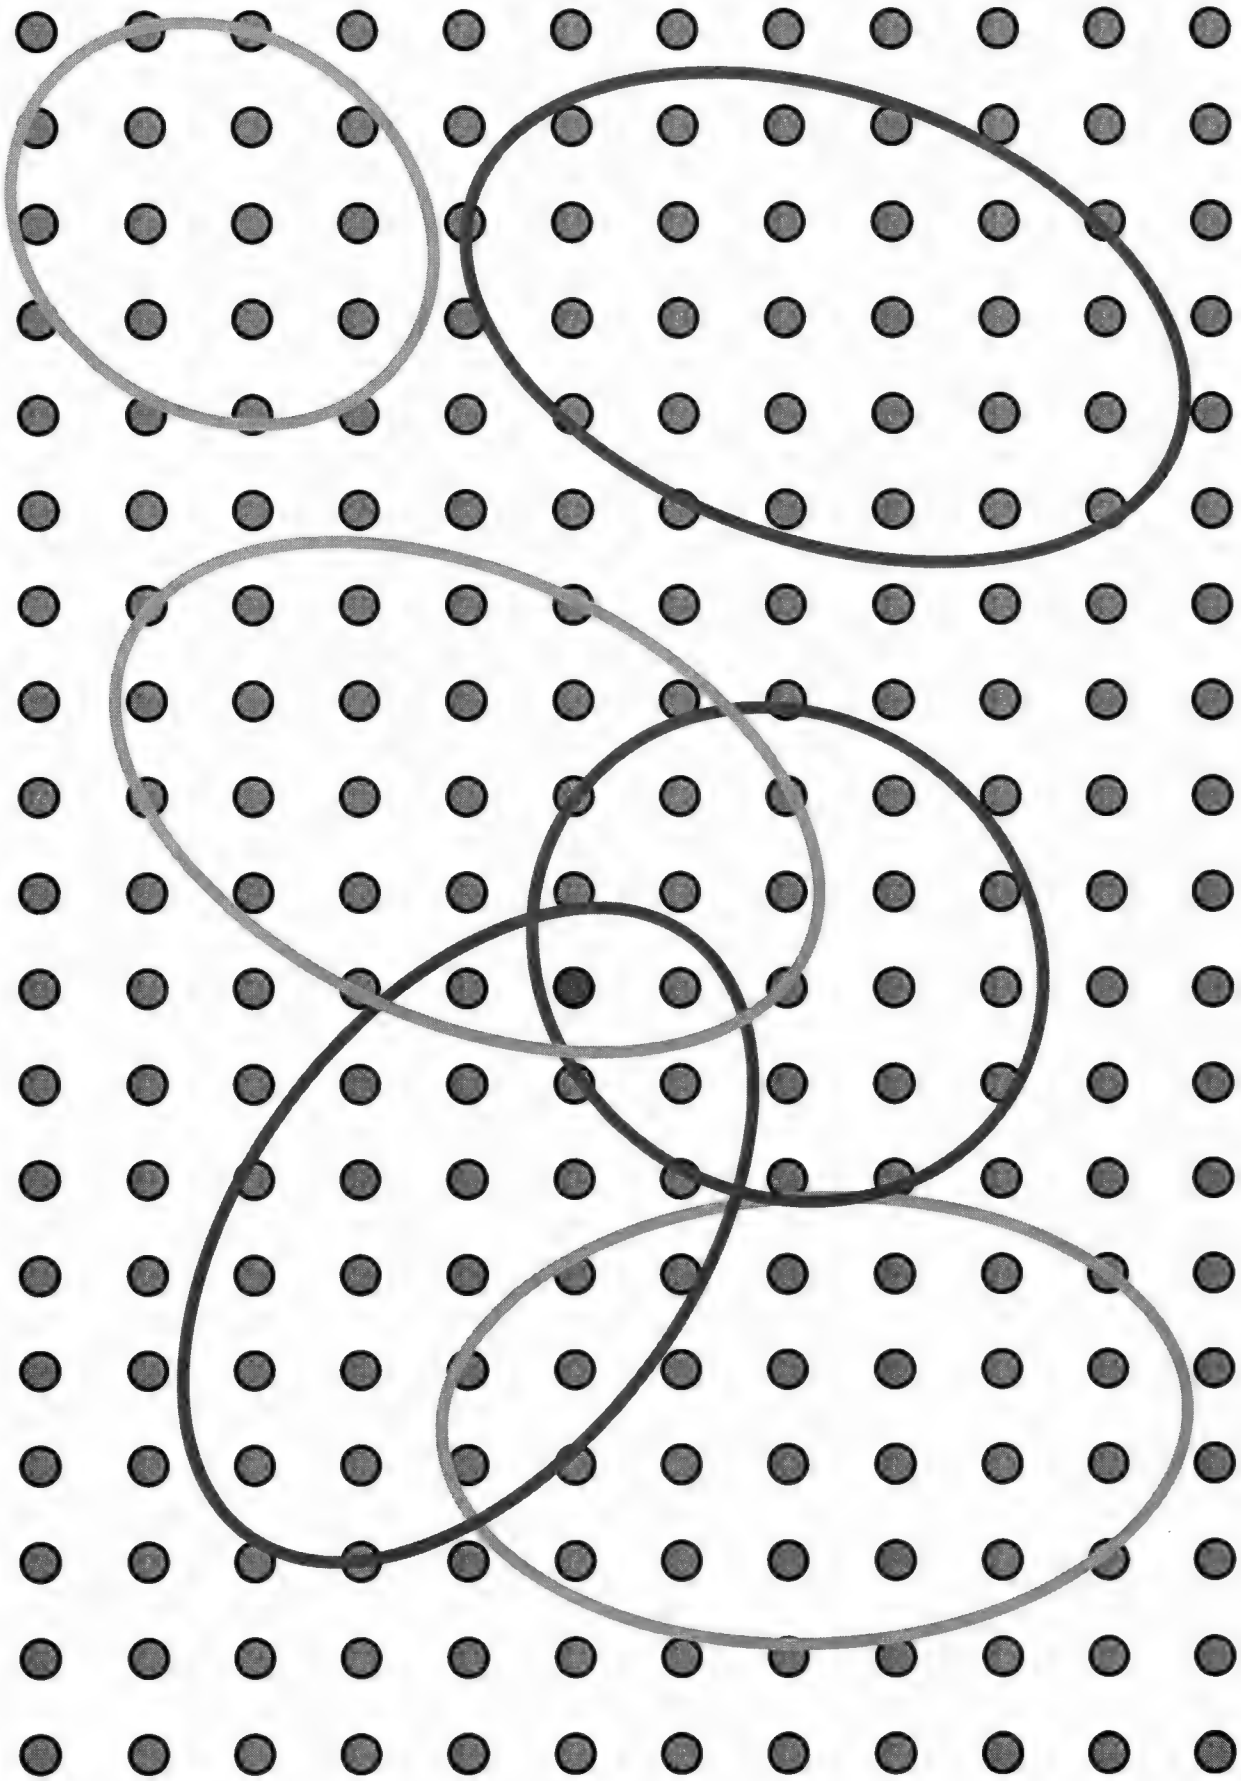

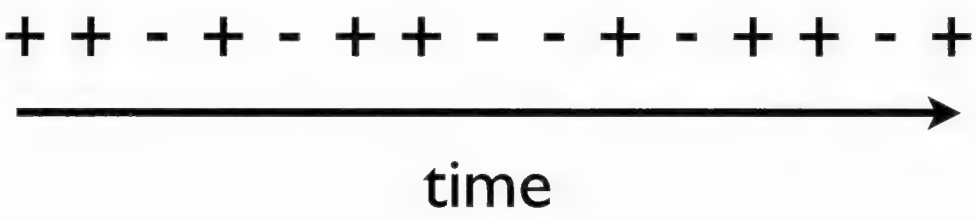
+ + - + - + + - - + - + + - +
time

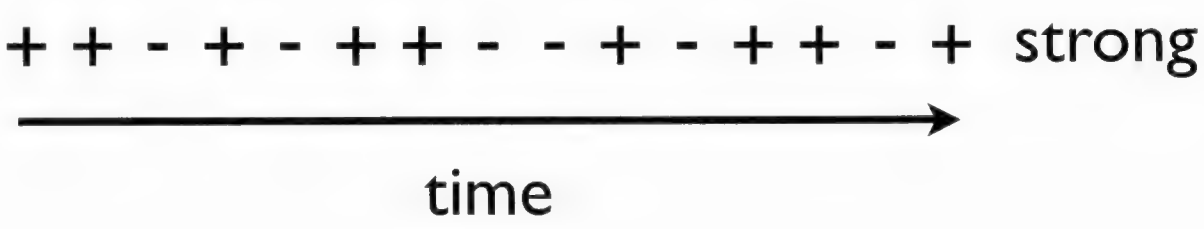
+ + - + - + + - - + - + + - +
strong
time

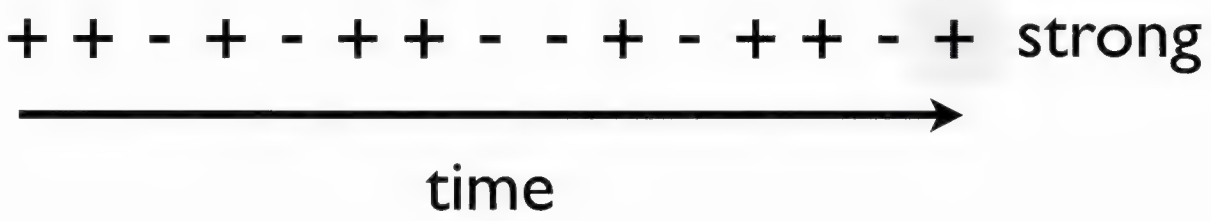
+ + - + - + + - - + - + + - +
strong
time

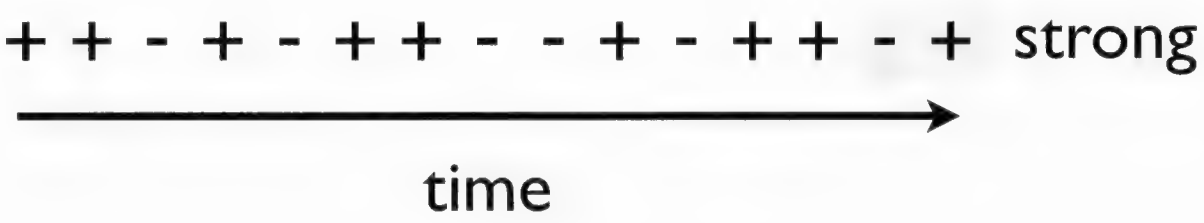
+ + - + - + + - - + - + + - +
strong
time

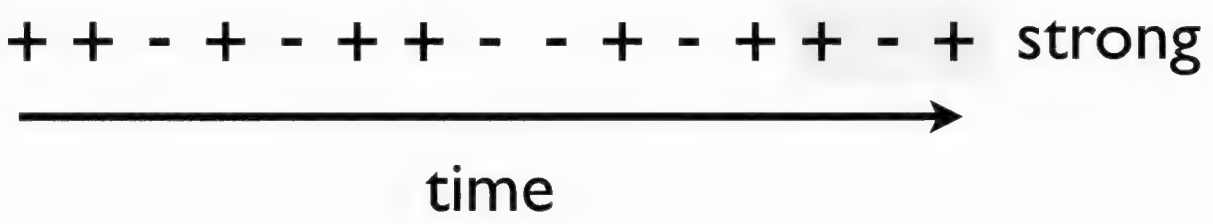
+ + - + - + + - - + - + + - +
strong
time

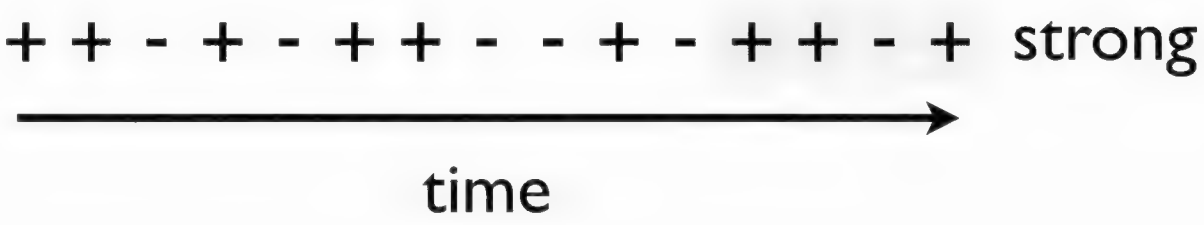
+ + - + - + + - - + - + + - +
strong
time

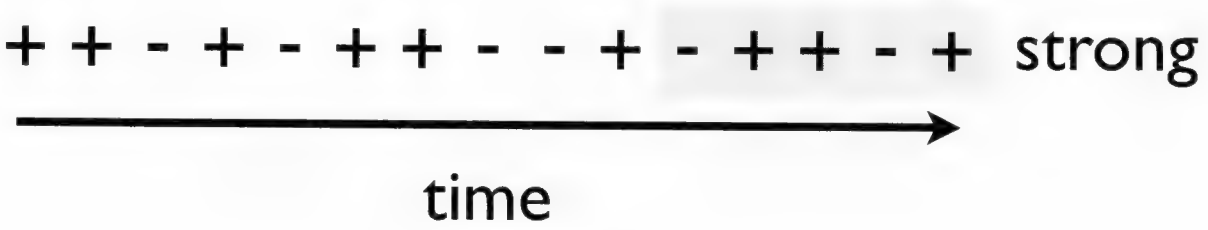
+ + - + - + + - - + - + + - +
strong
time

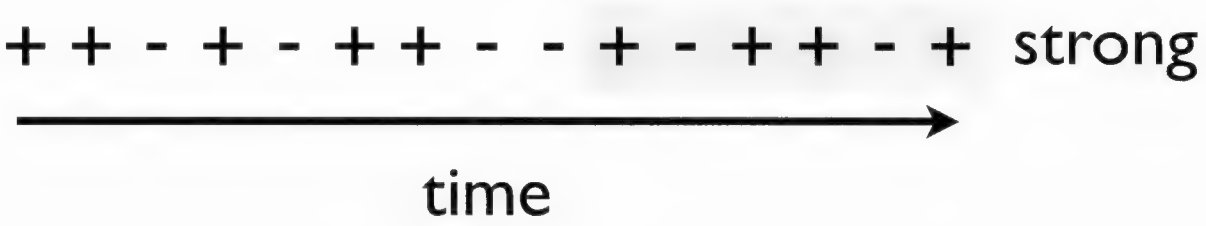
+ + - + - + + - - + - + + - +
strong
time

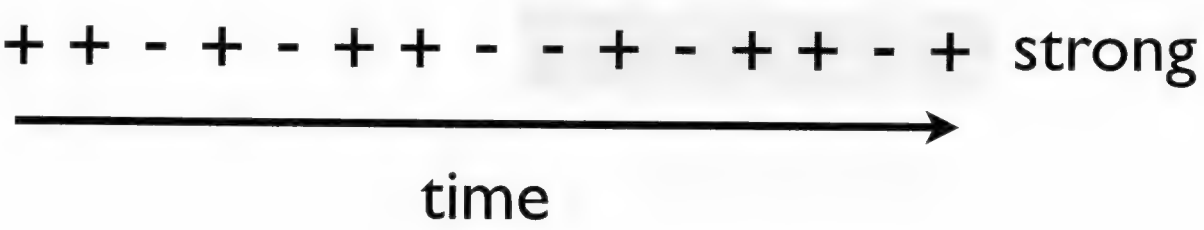
+ + - + - + + - - + - + + - +
strong
time

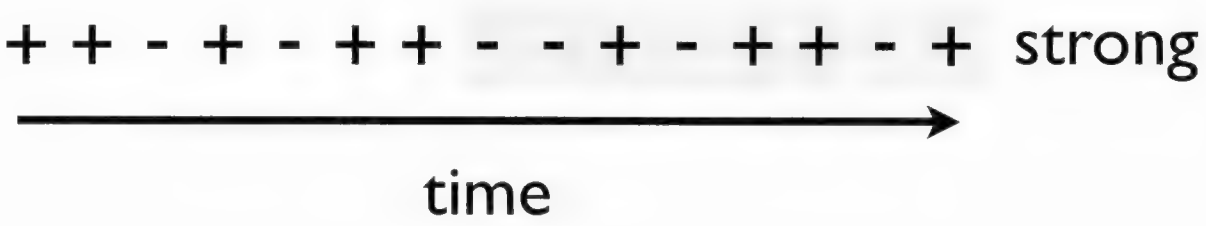
+ + - + - + + - - + - + + - +
strong
time

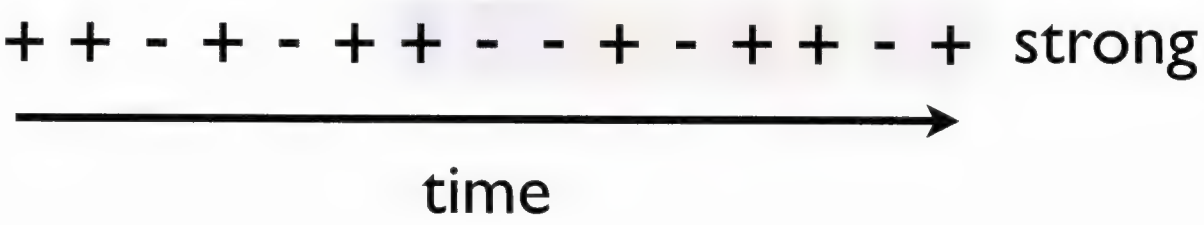
+ + - + - + + - - + - + + - +
strong
time

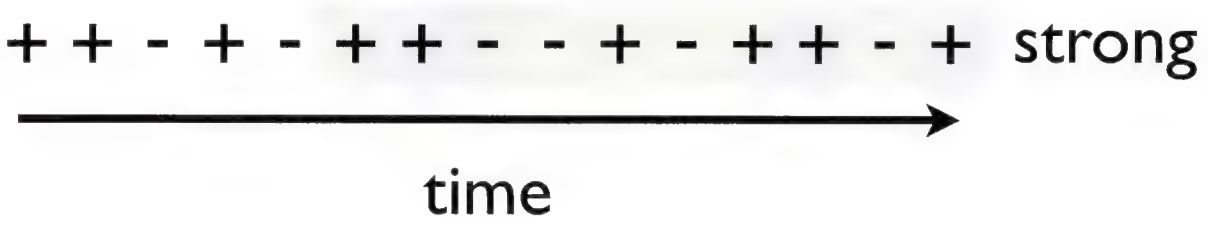
+ + - + - + + - - + - + + - +
strong
time

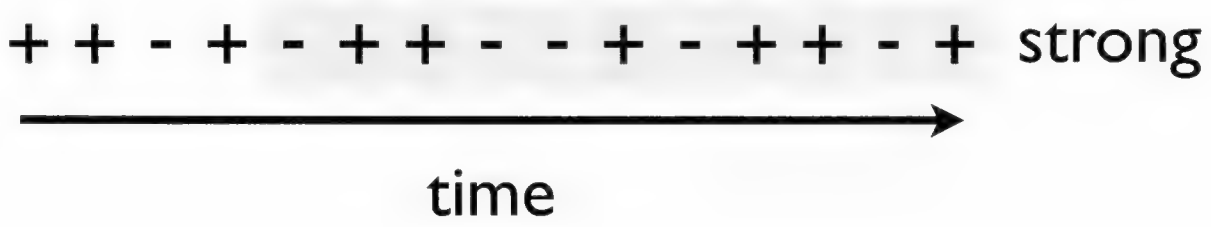
+ + - + - + + - - + - + + - +
strong
time

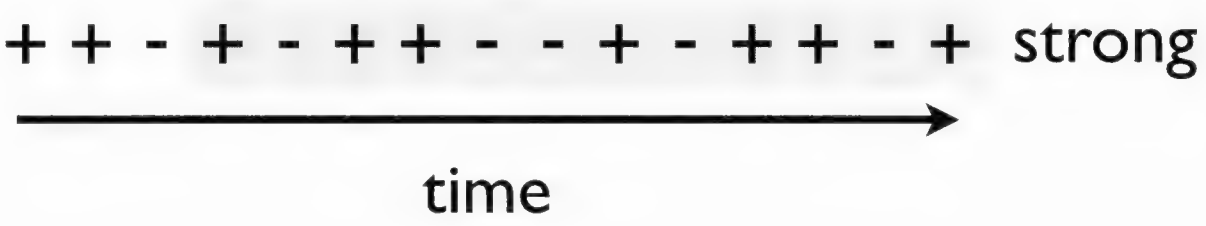
+ + - + - + + - - + - + + - +
strong
time

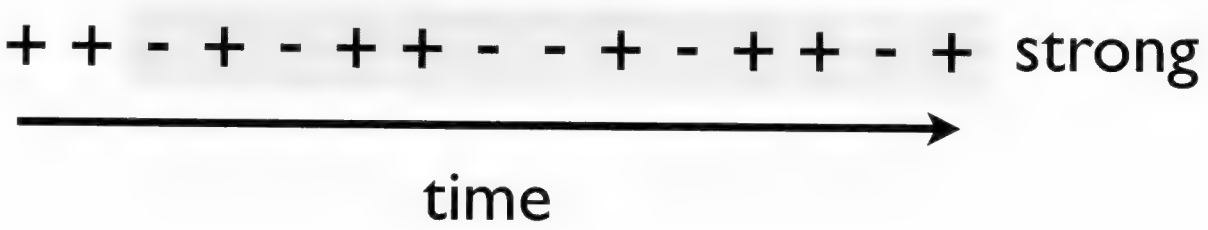
+ + - + - + + - - + - + + - +
strong
time

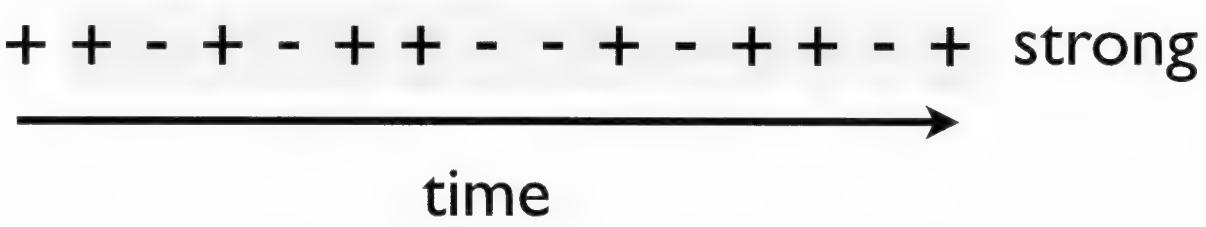
+ + - + - + + - - + - + + - +
strong
time

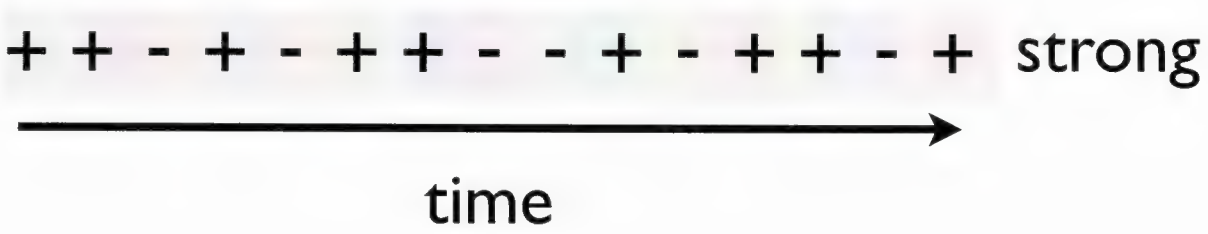
+ + - + - + + - - + - + + - +
strong
time

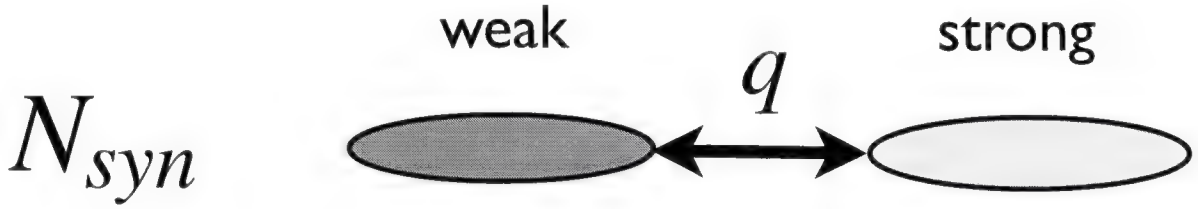
weak
strong
q
N_{syn}

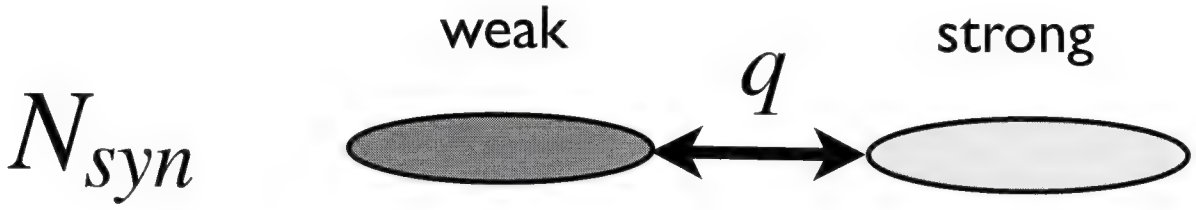

$$S = qN_{\mathrm{syn}} \exp(-qrt)$$

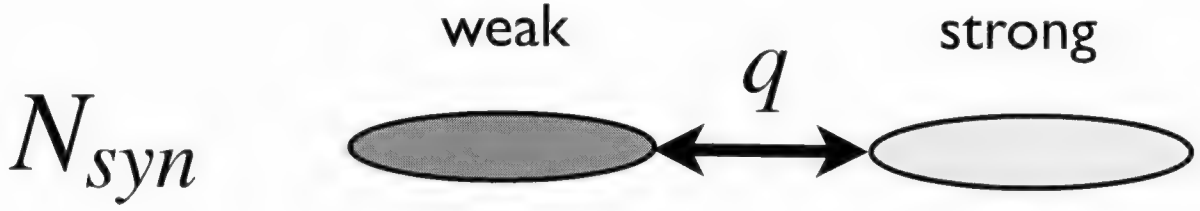

$$S = qN_{\mathrm{syn}} \exp(-qrt)$$

$$N = \sqrt{N_{\mathrm{syn}}}$$

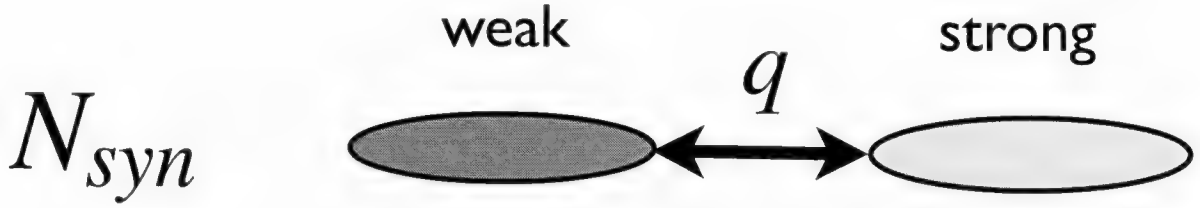

$$S = qN_{\mathrm{syn}} \exp(-qrt) \qquad N = \sqrt{N_{\mathrm{syn}}}$$

$$\frac{S}{N} = q\sqrt{N_{\mathrm{syn}}} \exp(-qrt)$$

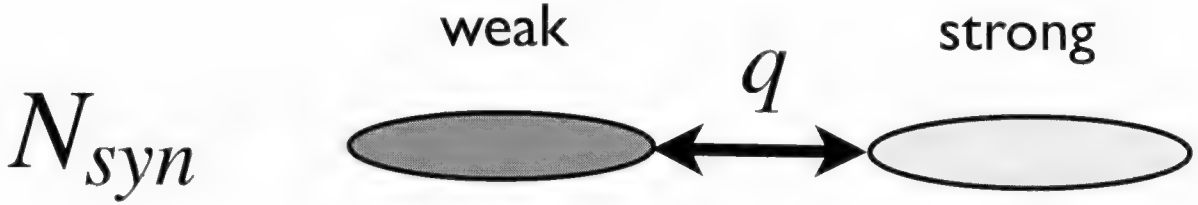

$$S = qN_{\mathrm{syn}} \exp(-qrt) \qquad N = \sqrt{N_{\mathrm{syn}}}$$

$$\frac{S}{N} = q\sqrt{N_{\mathrm{syn}}} \exp(-qrt)$$

$$\frac{S_0}{N_0} = q\sqrt{N_{\mathrm{syn}}}$$

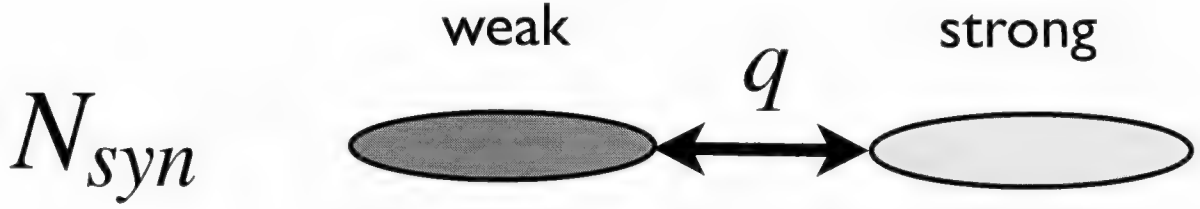

$$S = qN_{\text{syn}} \exp(-qrt) \qquad N = \sqrt{N_{\text{syn}}}$$

$$\frac{S}{N} = q\sqrt{N_{\text{syn}}} \exp(-qrt)$$

$$\frac{S_0}{N_0} = q\sqrt{N_{\text{syn}}} \qquad t_{\max} = \frac{\ln(q\sqrt{N_{\text{syn}}})}{qr}$$

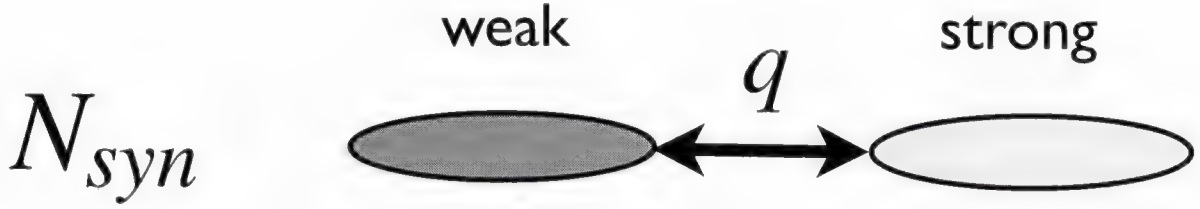

$$S = qN_{\text{syn}} \exp(-qrt) \qquad N = \sqrt{N_{\text{syn}}}$$

$$\frac{S}{N} = q\sqrt{N_{\text{syn}}} \exp(-qrt)$$

$$\boxed{\frac{S_0}{N_0} = q\sqrt{N_{\text{syn}}}} \qquad \boxed{t_{\max} = \frac{\ln(q\sqrt{N_{\text{syn}}})}{qr}}$$

Amit & Fusi (1992)
Fusi (2002)

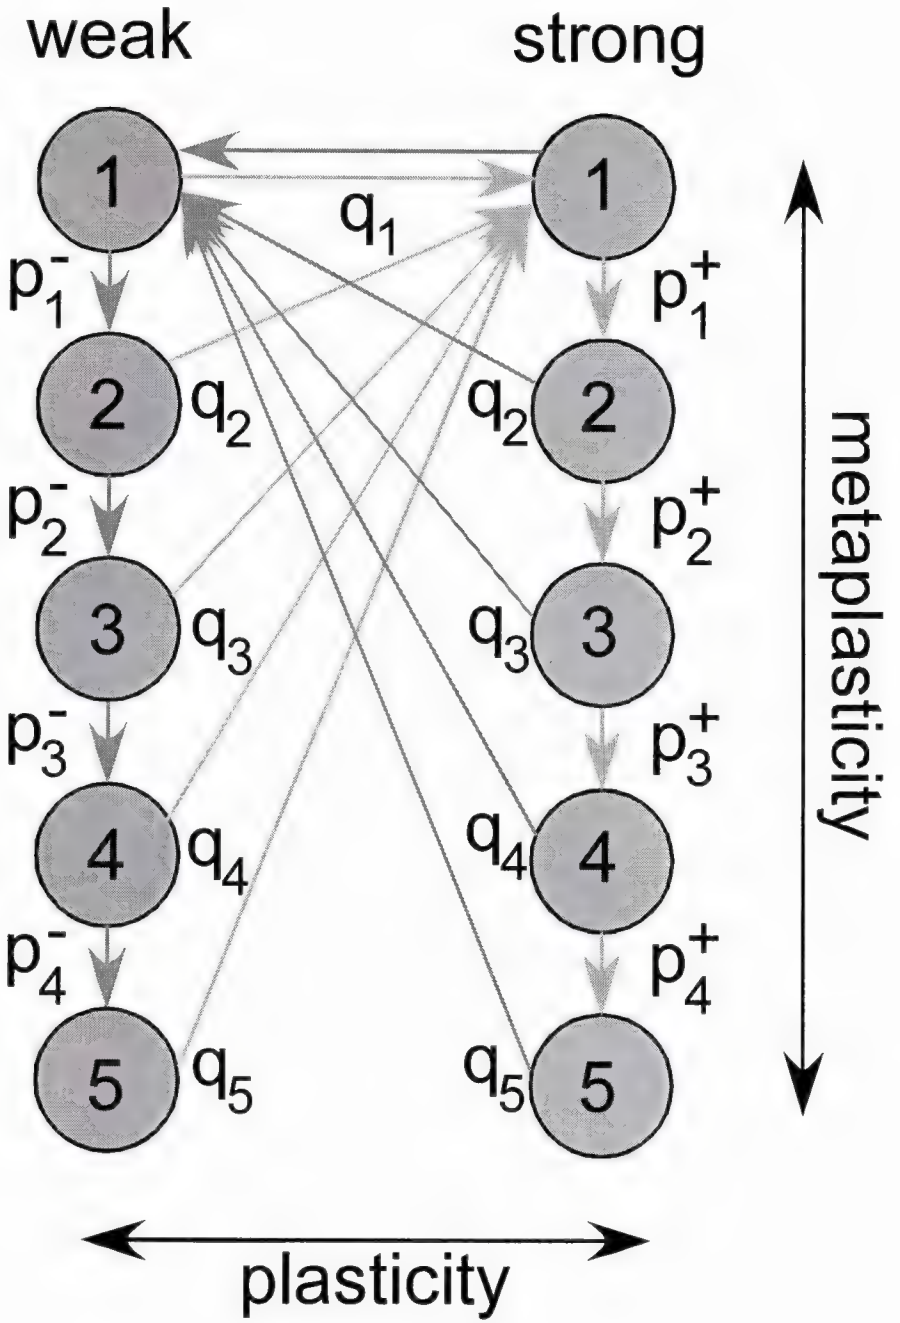
weak
strong
1
1
q_1
p_1^-
p_1^+
2
q_2
q_2
2
p_2^-
p_2^+
3
q_3
q_3
3
p_3^-
p_3^+
4
q_4
q_4
4
p_4^-
p_4^+
5
q_5
q_5
5
metaplasticity
plasticity

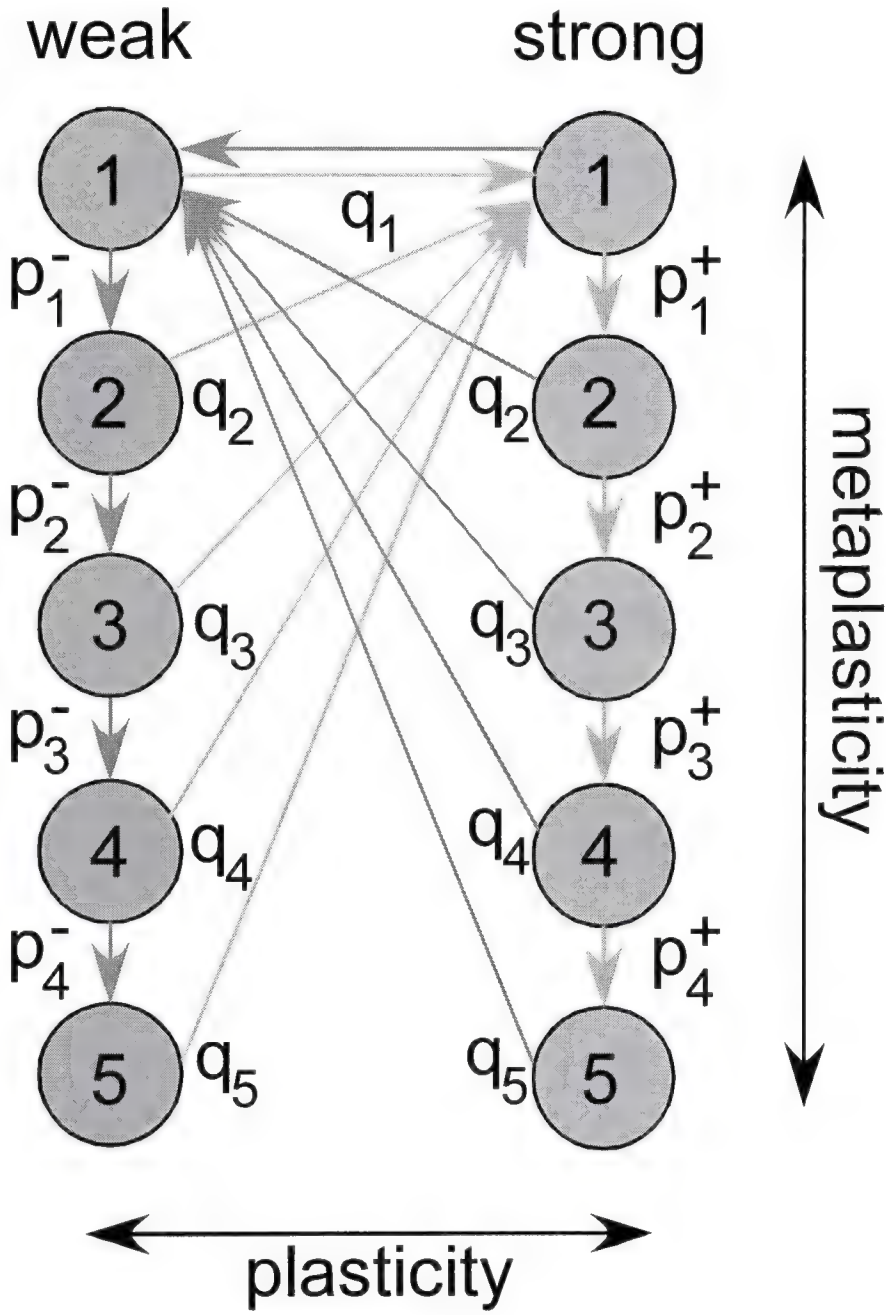

with Stefano Fusi and Patrick Drew

cascade

$$\frac{S_0}{N_0} \propto \frac{\sqrt{N_{\mathrm{syn}}}}{n} \qquad t_{\max} \propto 2^n = \sqrt{N_{\mathrm{syn}}}$$

cascade

$$\frac{S_0}{N_0} \propto \frac{\sqrt{N_{\mathrm{syn}}}}{n} \qquad t_{\max} \propto 2^n = \sqrt{N_{\mathrm{syn}}}$$

$$\frac{S_0}{N_0} \propto \frac{\sqrt{N_{\mathrm{syn}}}}{\ln\left(\sqrt{N_{\mathrm{syn}}}\right)} \qquad t_{\max} \propto \sqrt{N_{\mathrm{syn}}}$$

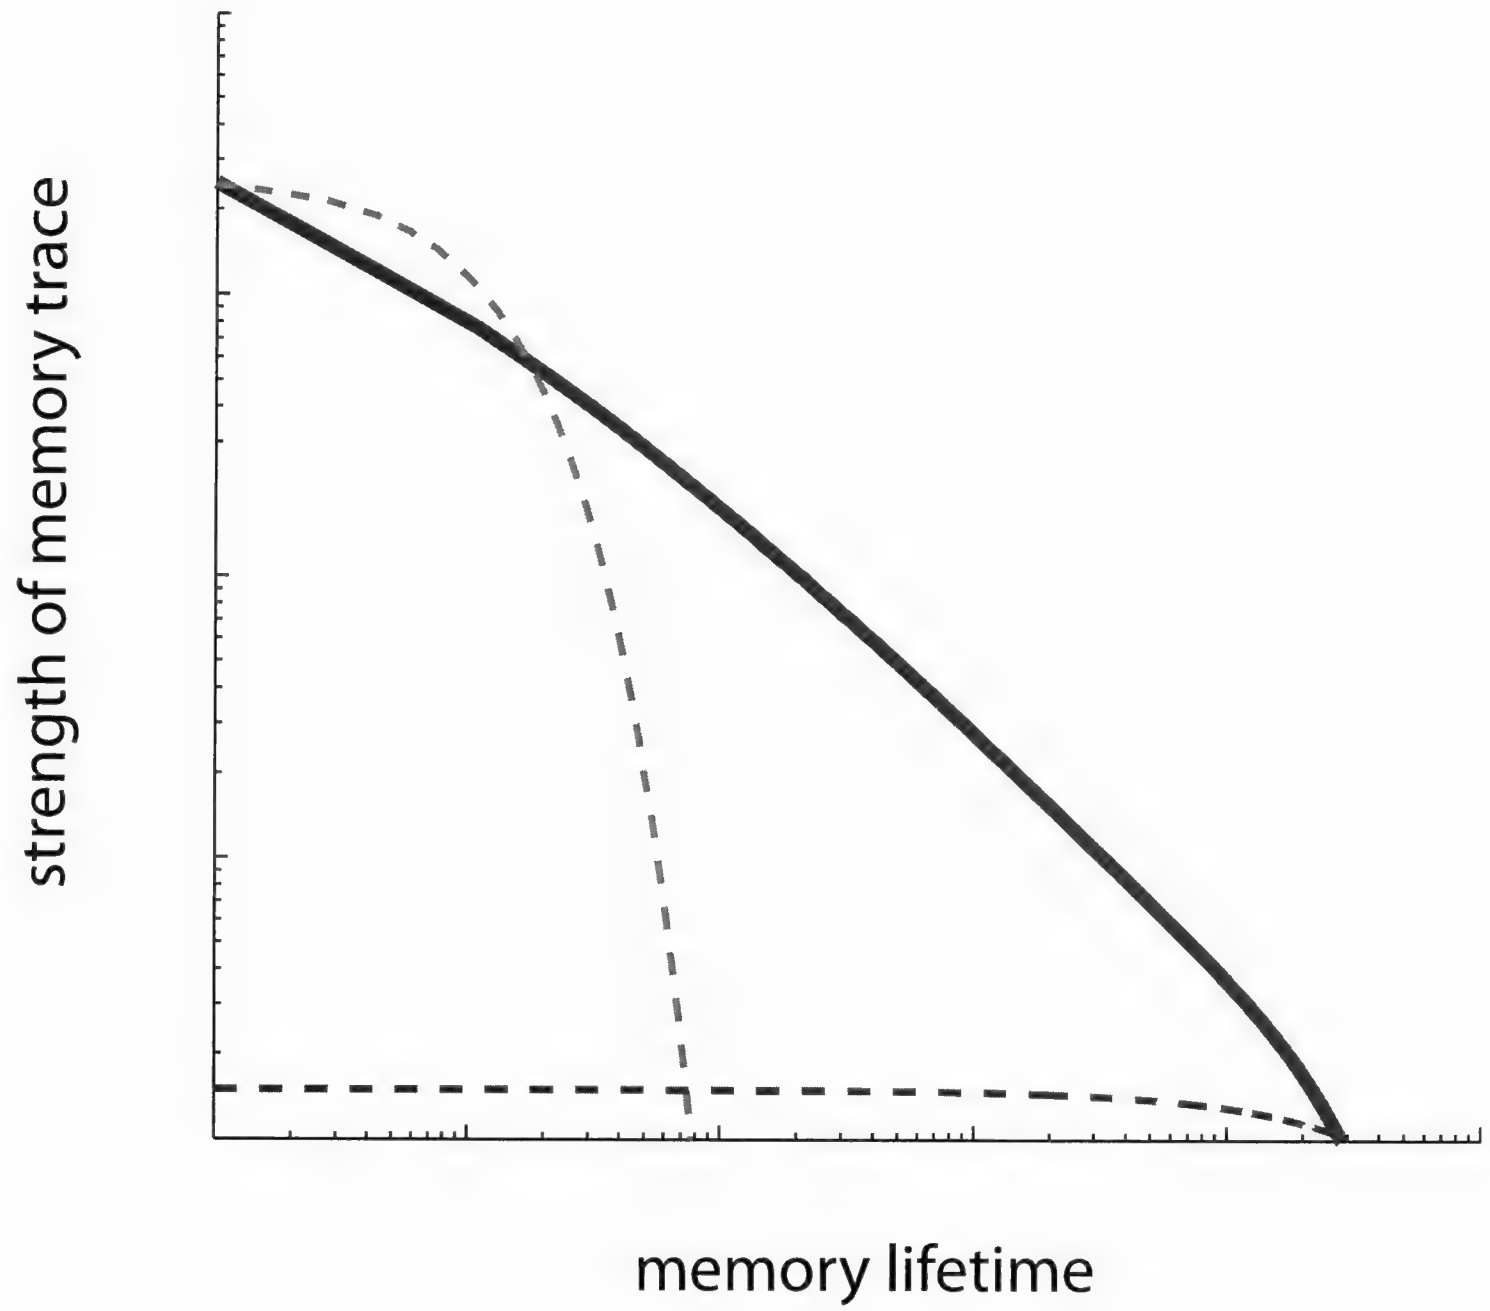
strength of memory trace
memory lifetime

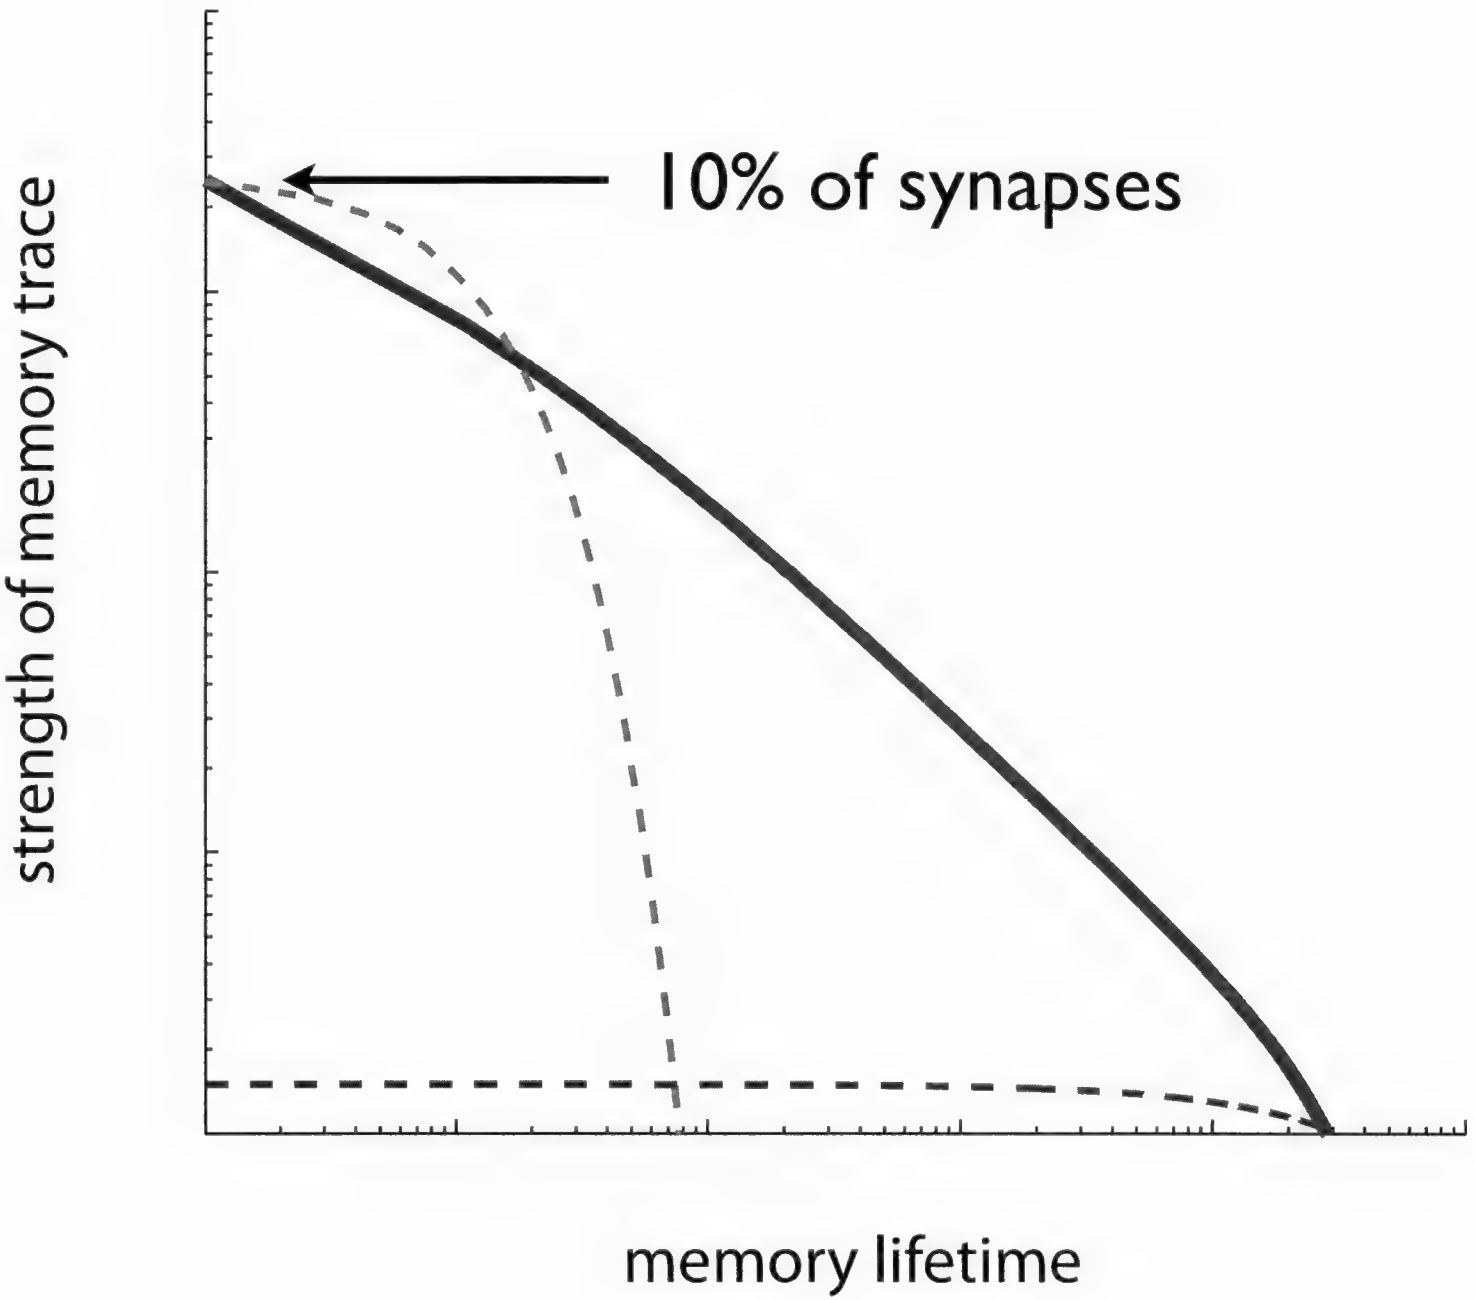
10% of synapses
strength of memory trace
memory lifetime

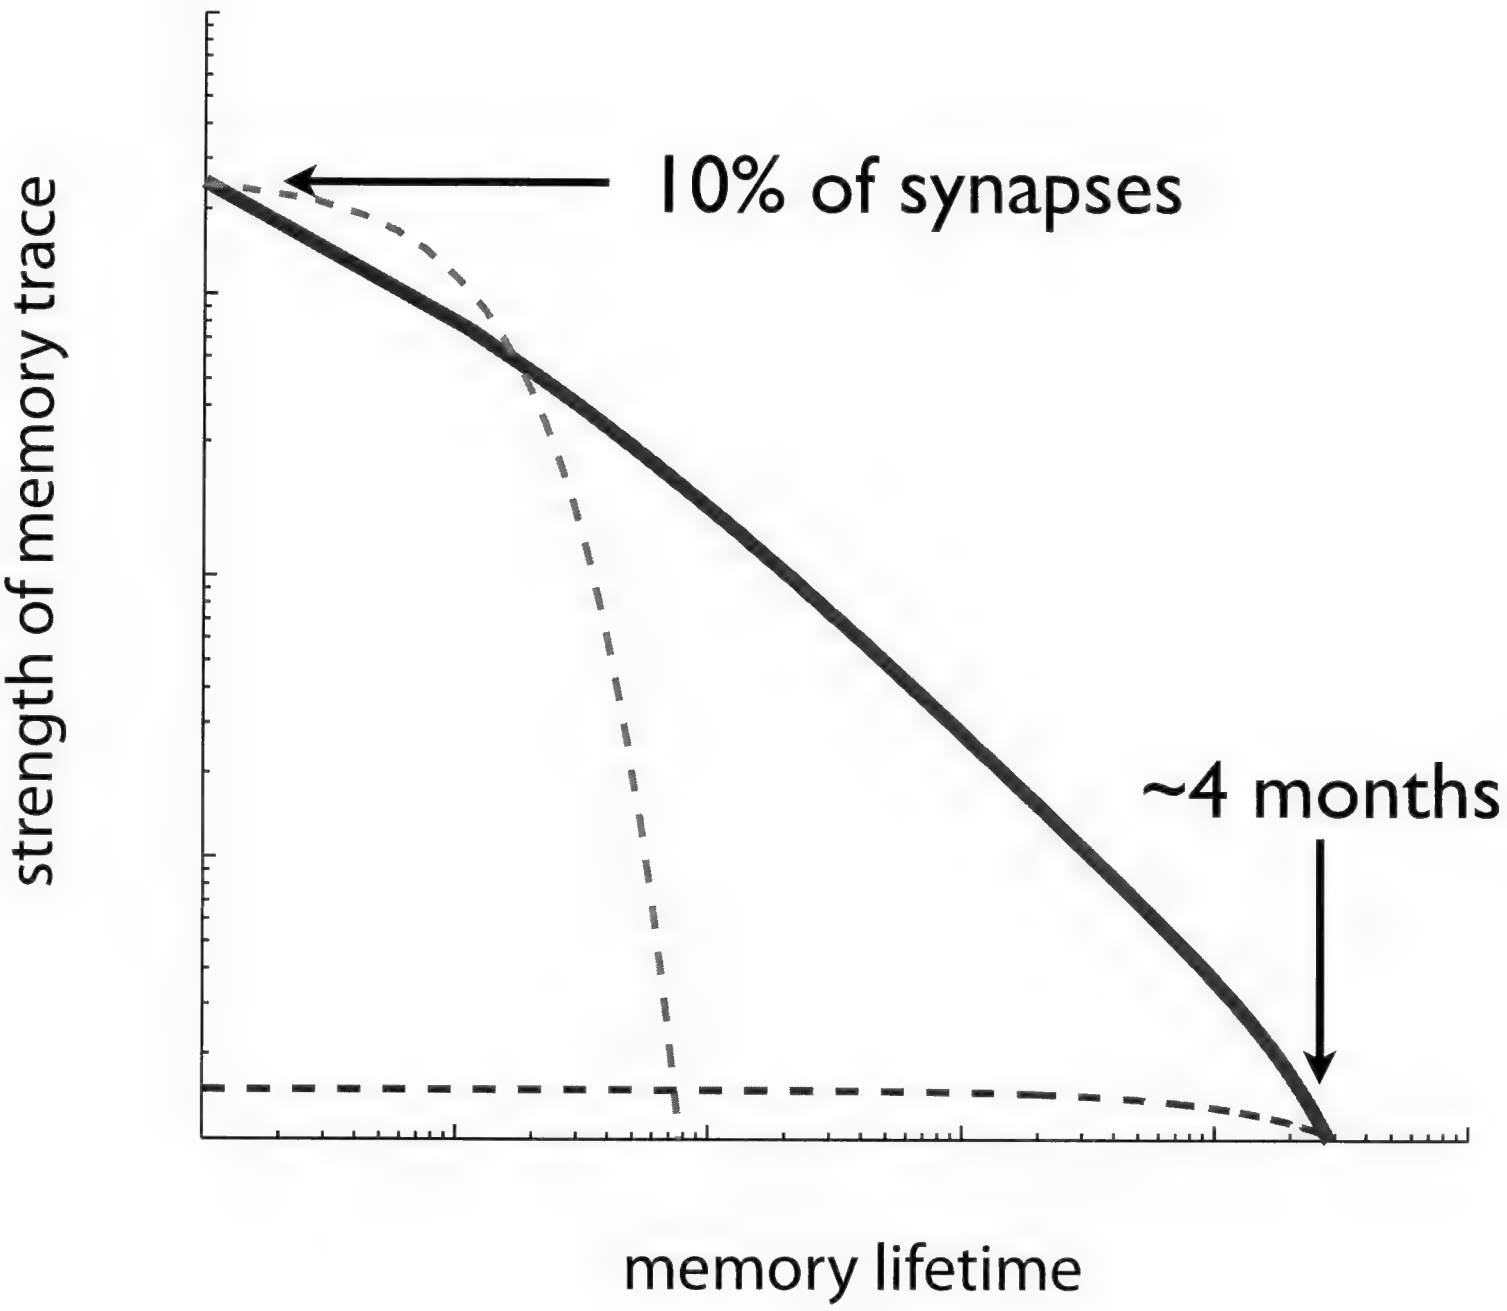
10% of synapses
~4 months
strength of memory trace
memory lifetime

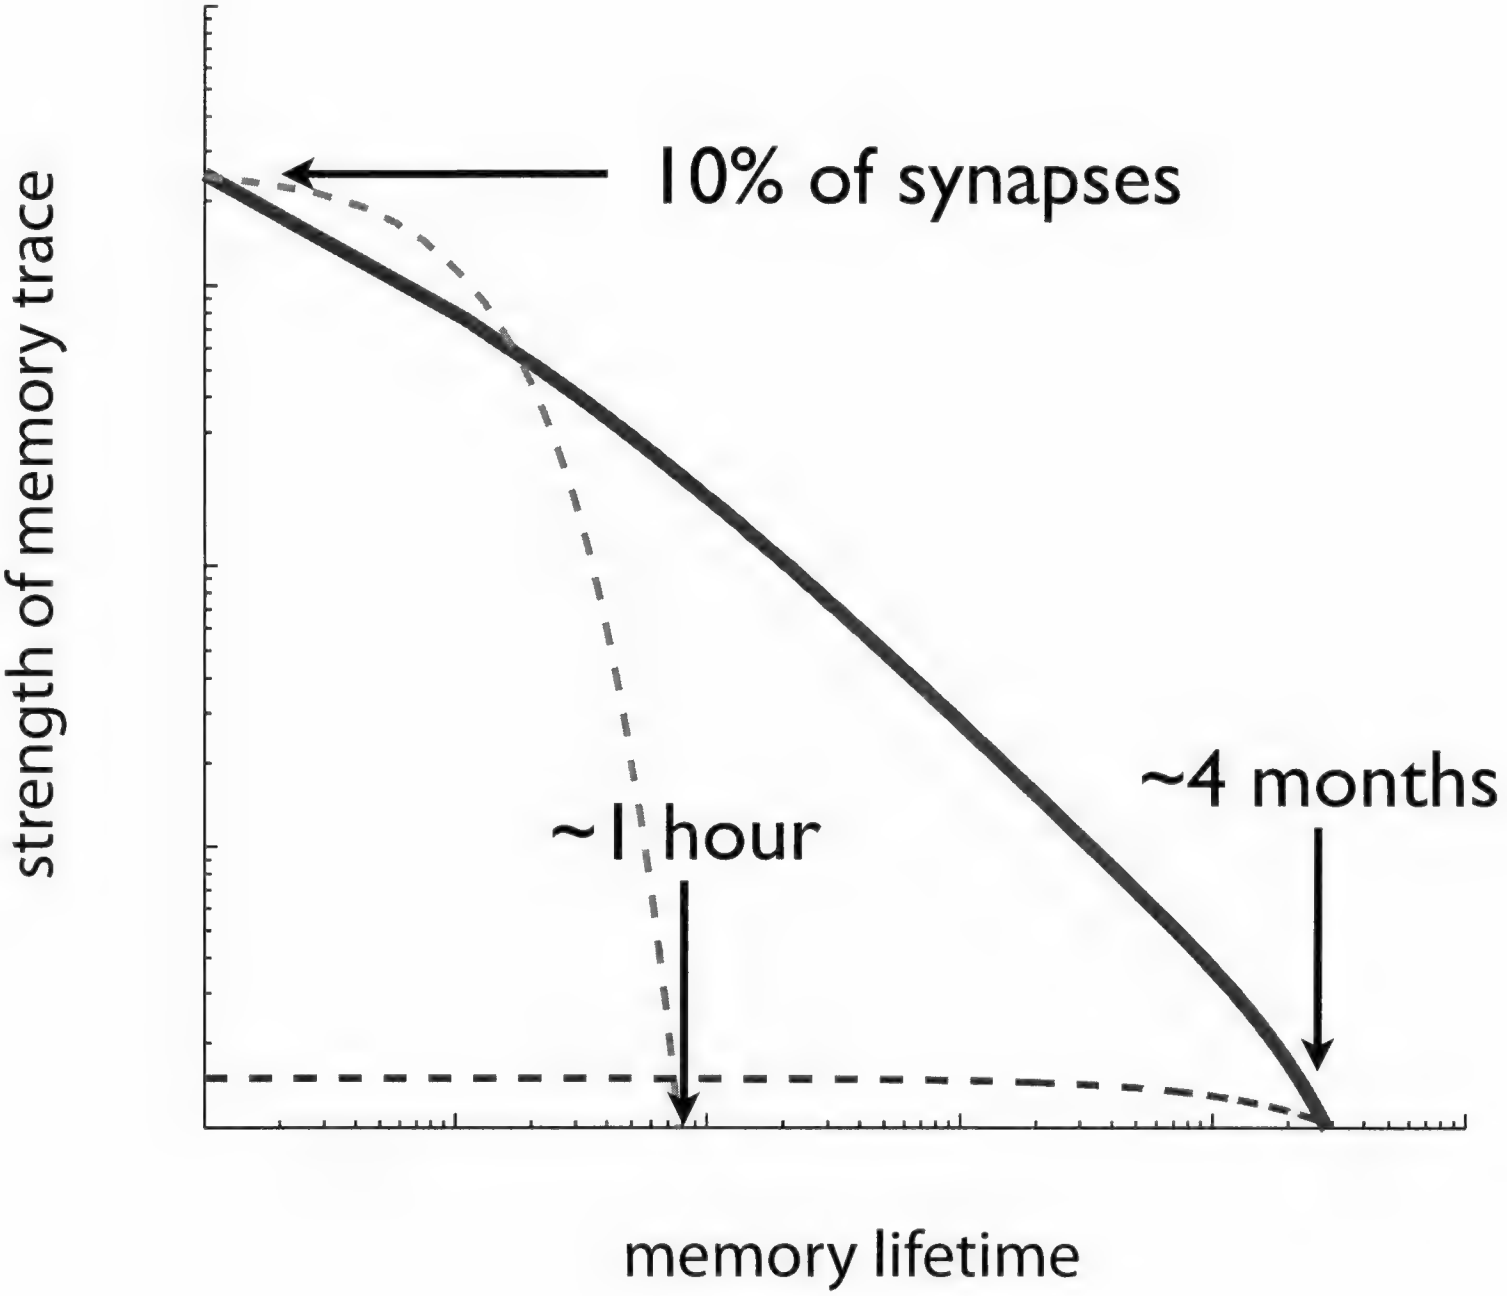
10% of synapses
~4 months
~1 hour
strength of memory trace
memory lifetime

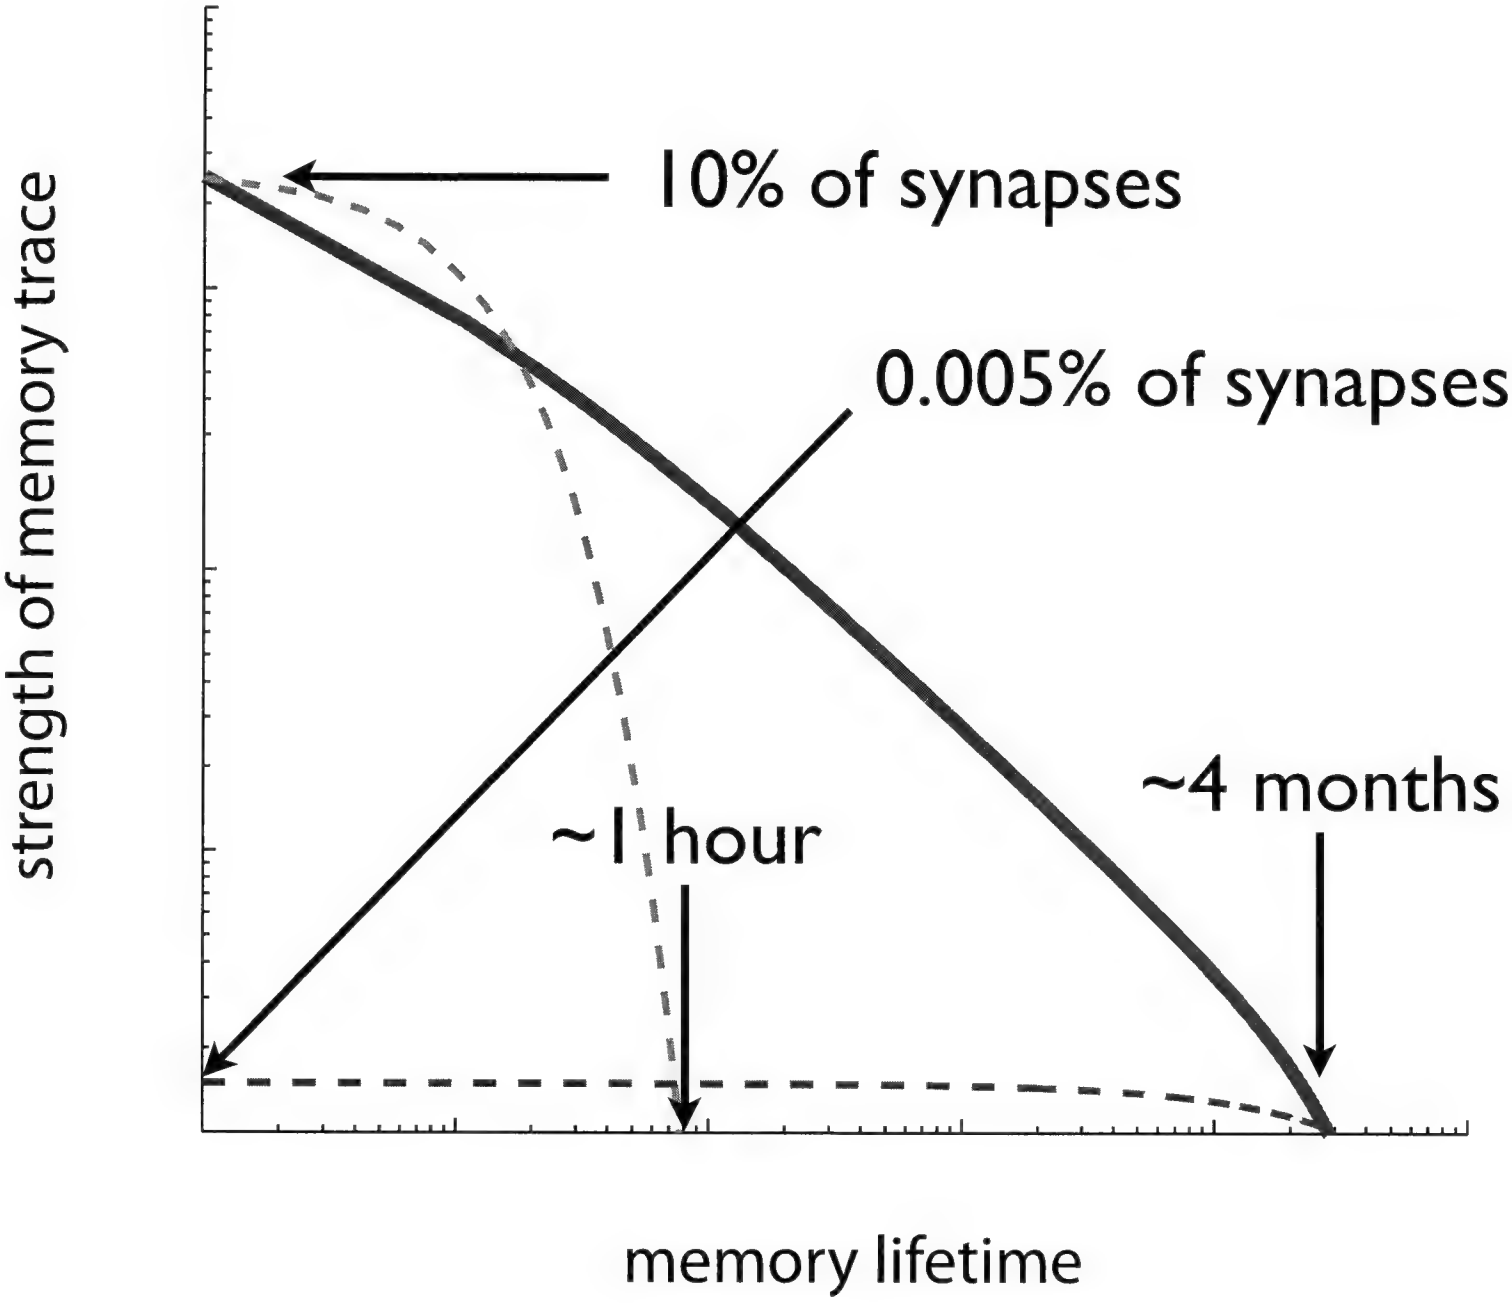
10% of synapses
0.005% of synapses
~4 months
~1 hour
strength of memory trace
memory lifetime

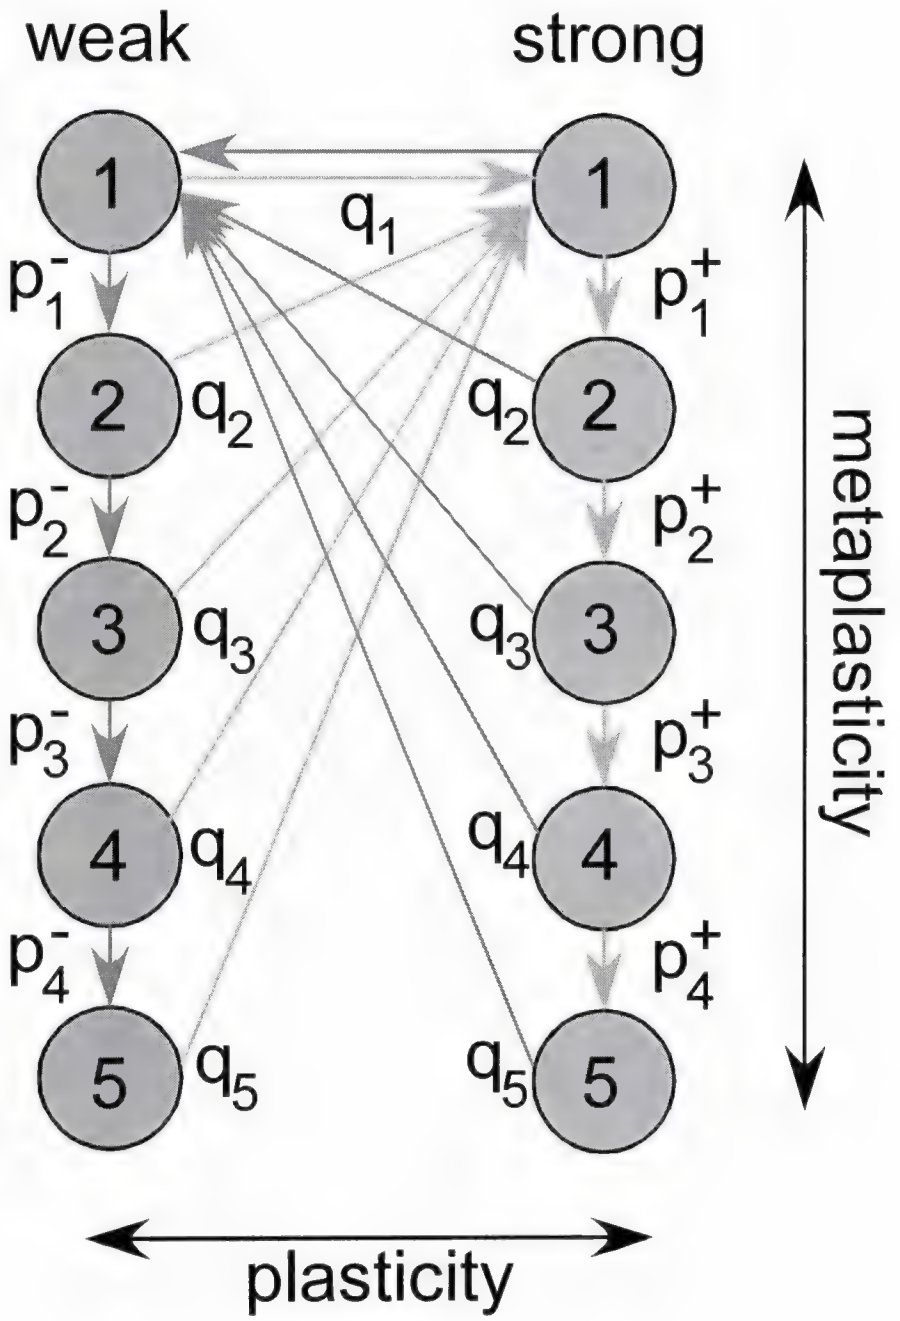
weak
strong
1
1
q_1
p_1^-
p_1^+
2
q_2
q_2
2
p_2^-
p_2^+
3
q_3
q_3
3
p_3^-
p_3^+
4
q_4
q_4
4
p_4^-
p_4^+
5
q_5
q_5
5
metaplasticity
plasticity

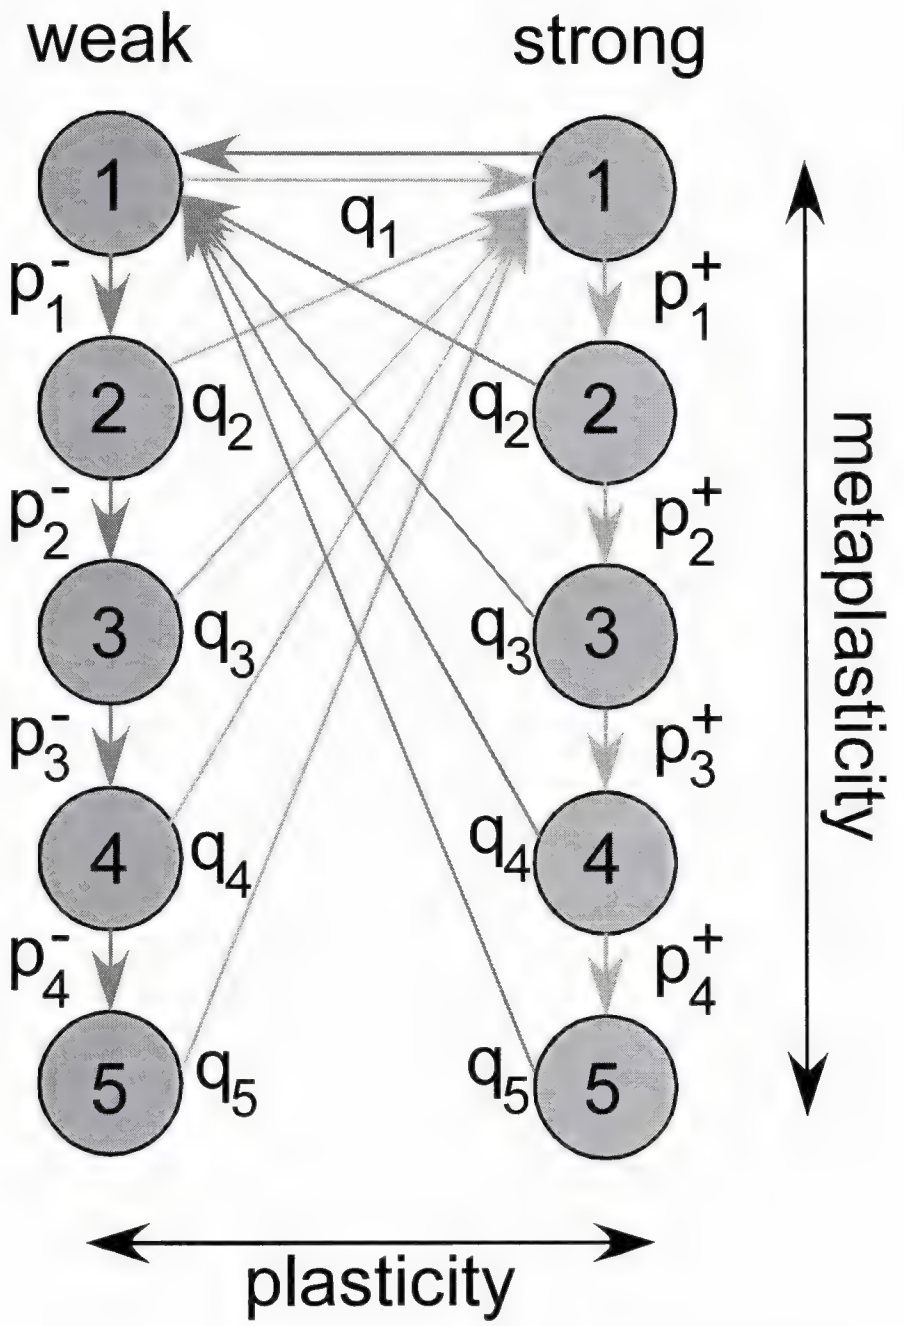

high plasticity for learning
low plasticity for retention

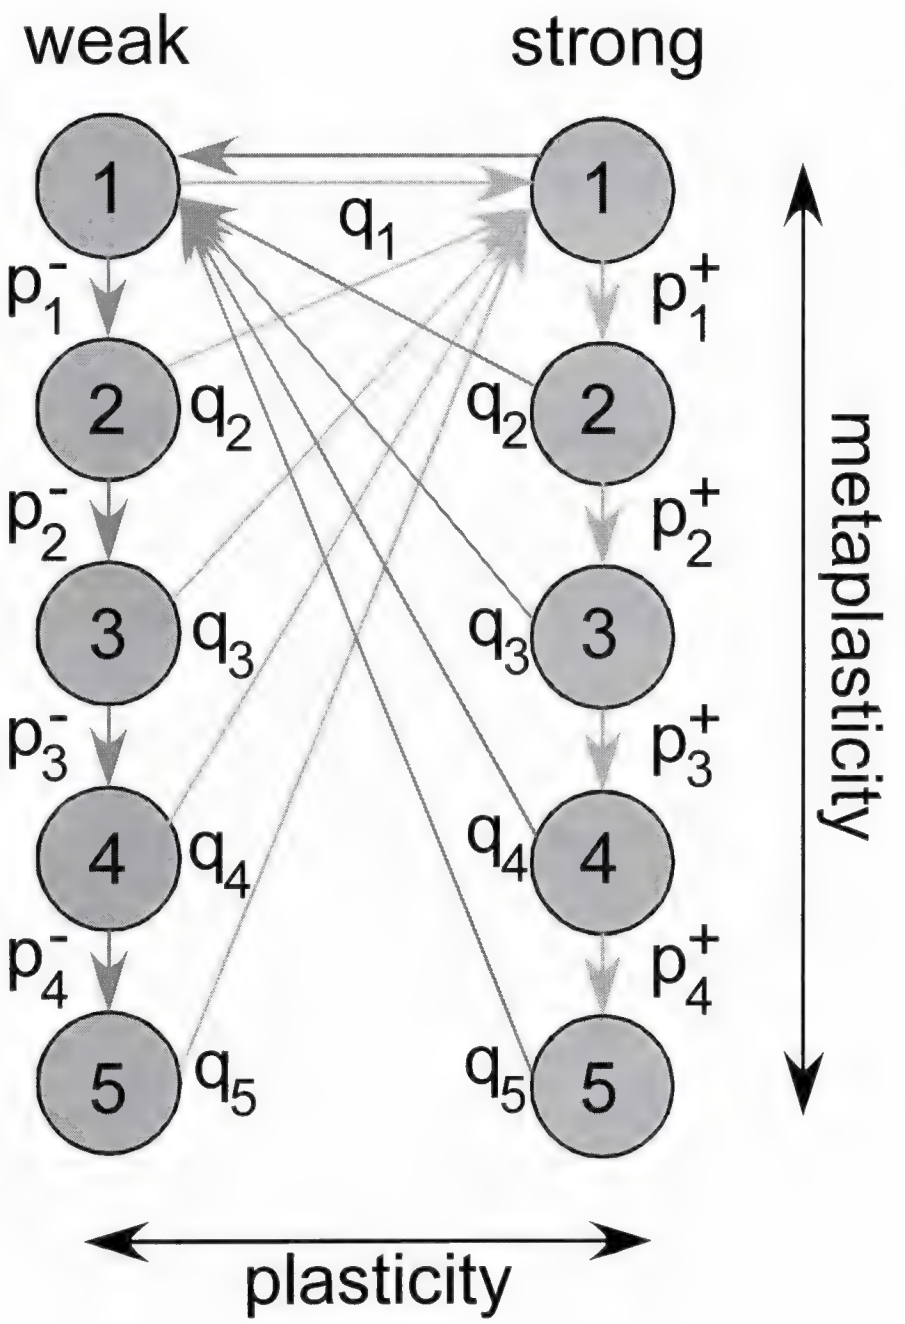

high plasticity for learning
low plasticity for retention

single-trial learning

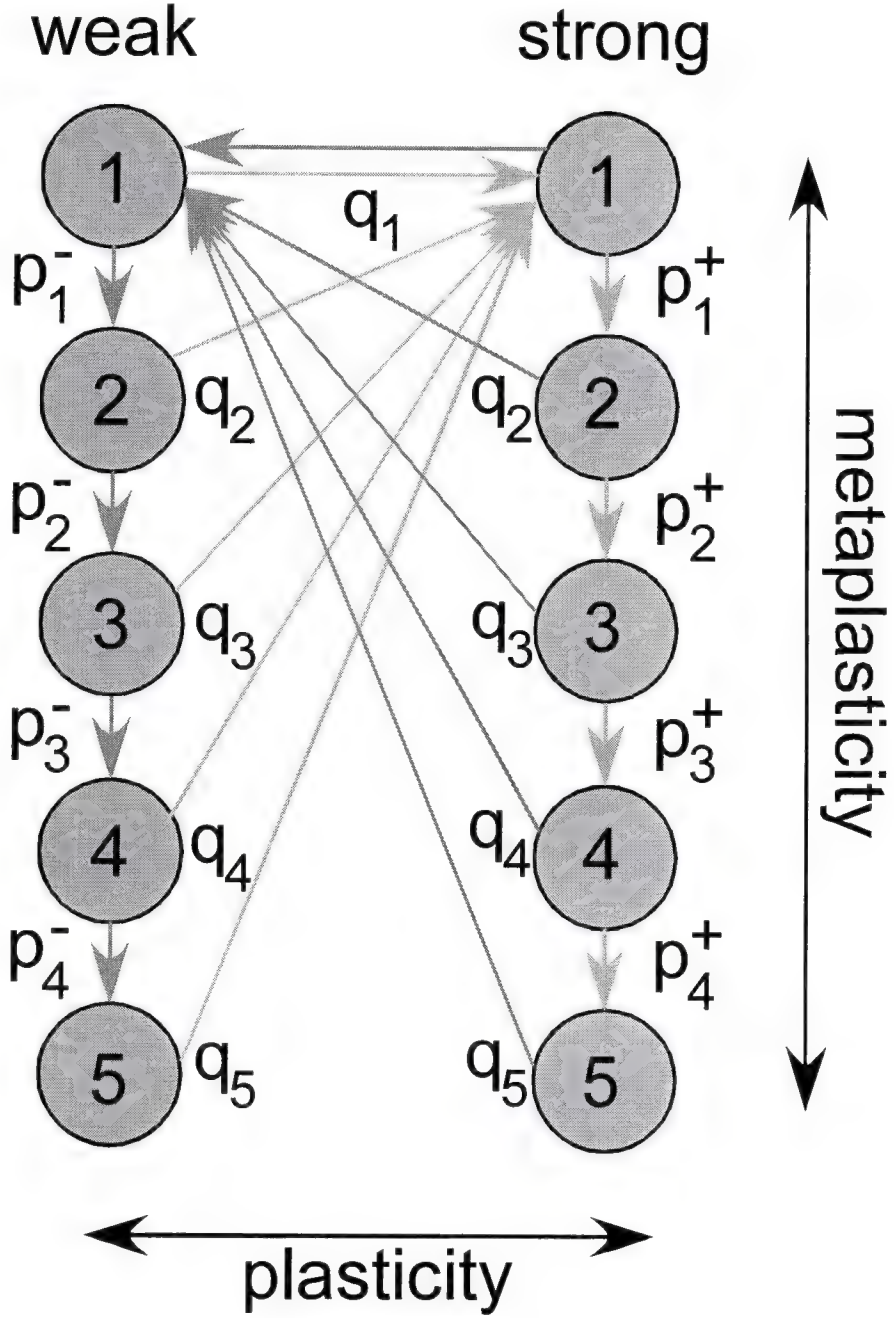

high plasticity for learning
low plasticity for retention

single-trial learning

spaced/massed effect

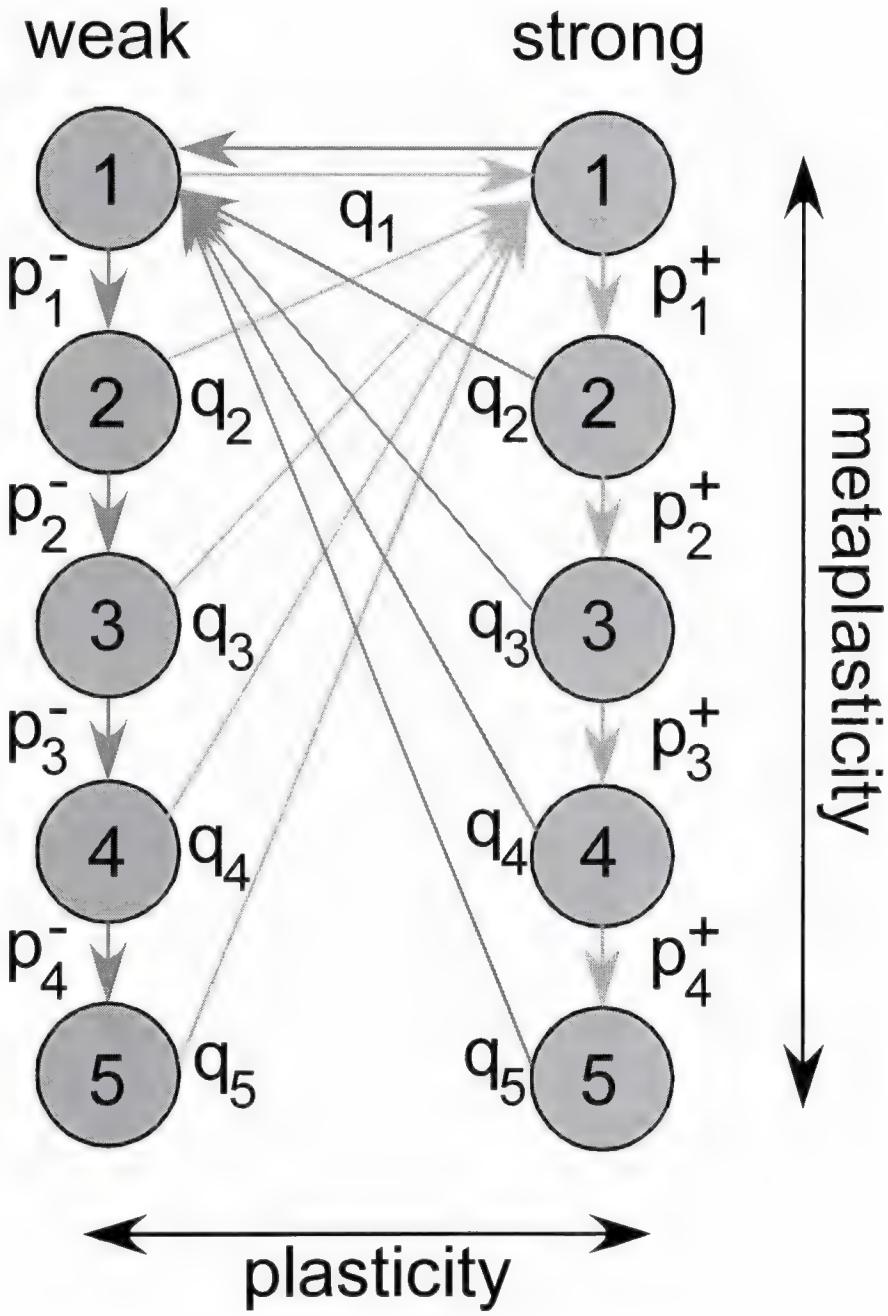

high plasticity for learning
low plasticity for retention

single-trial learning

spaced/massed effect

reinforcement learning

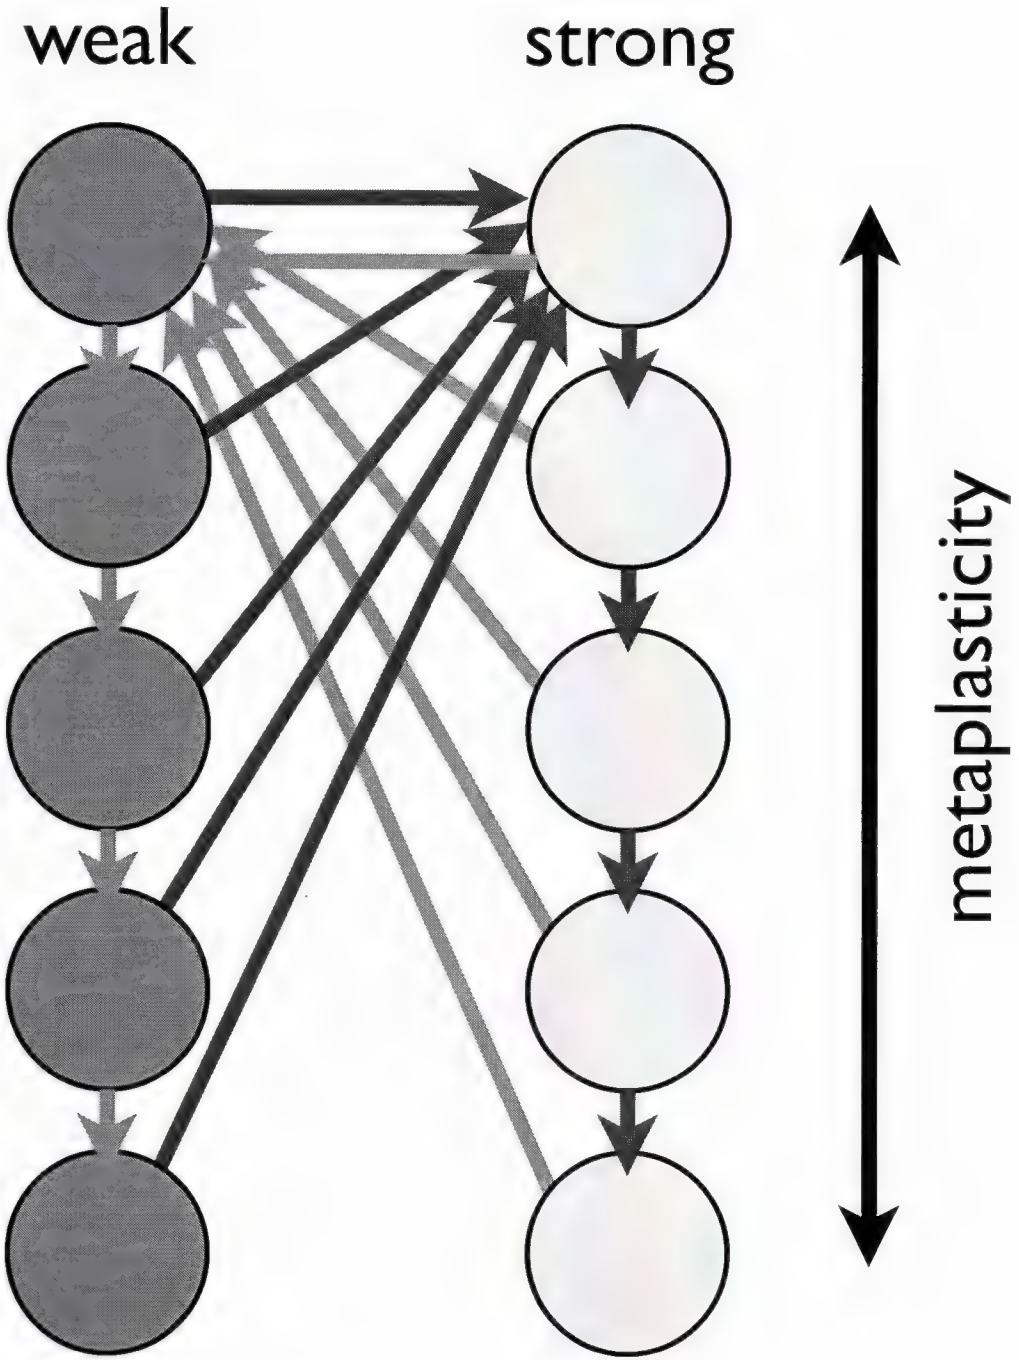
weak
strong
metaplasticity

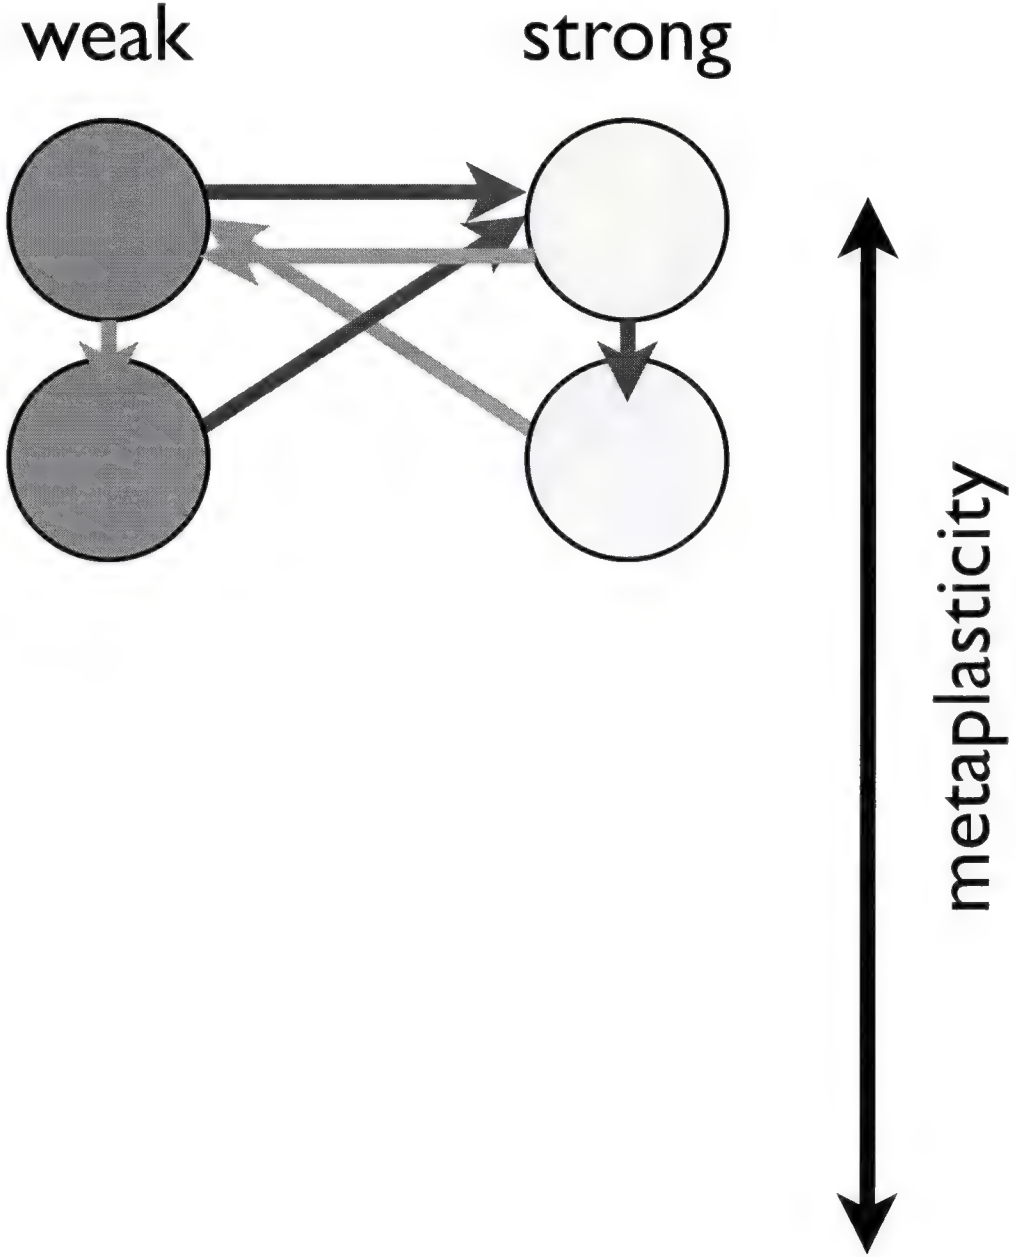
weak
strong
metaplasticity

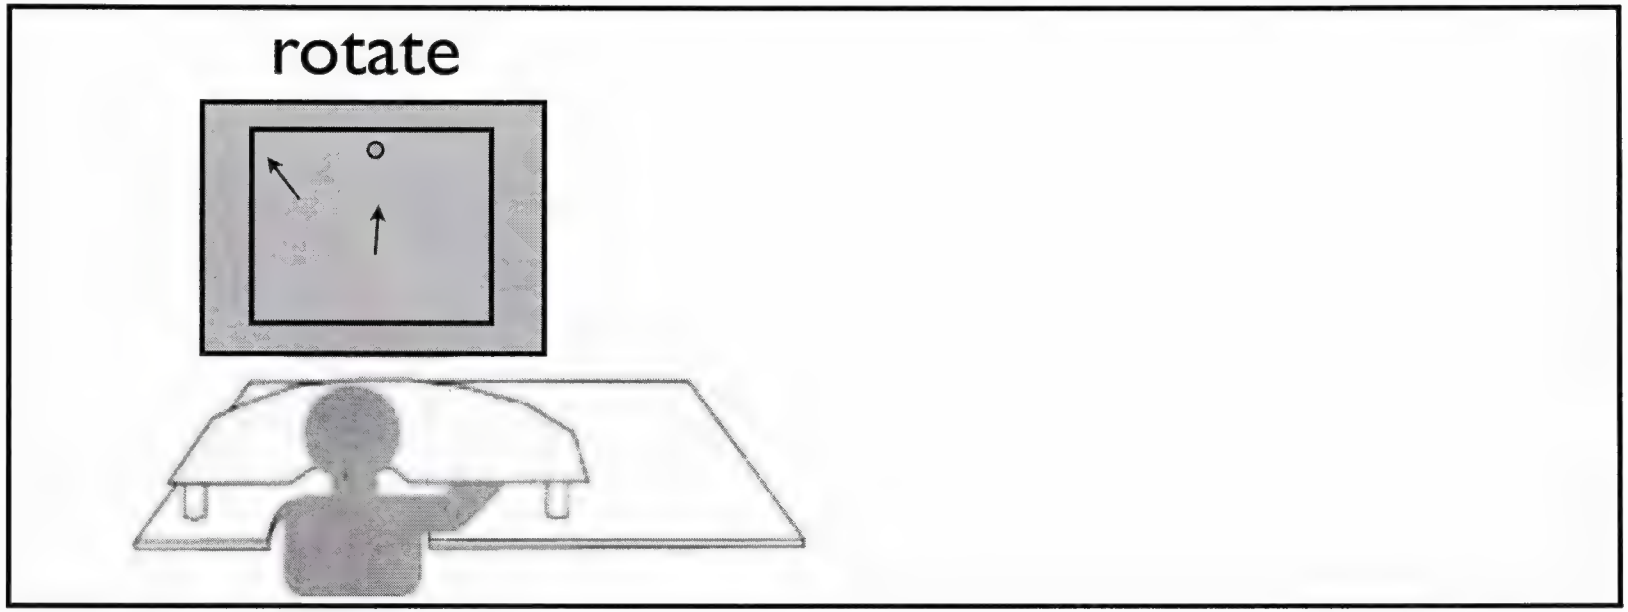

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

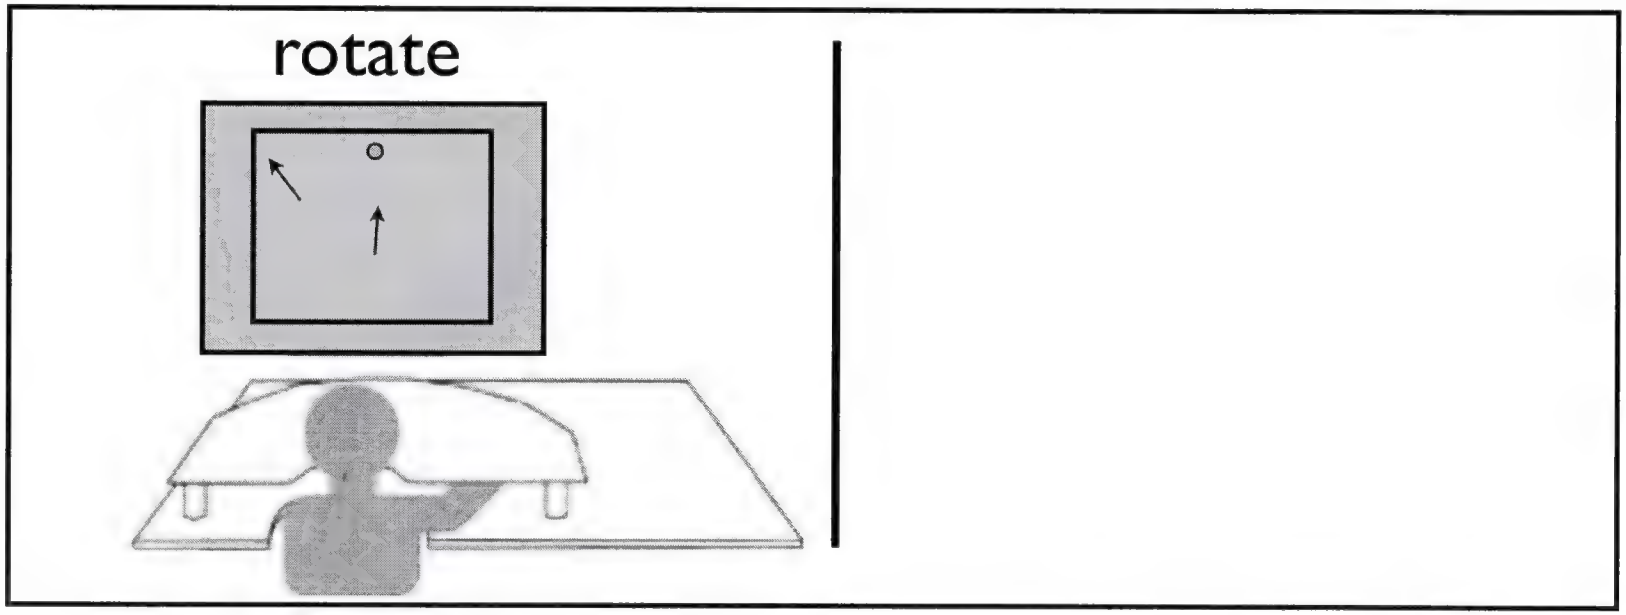

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

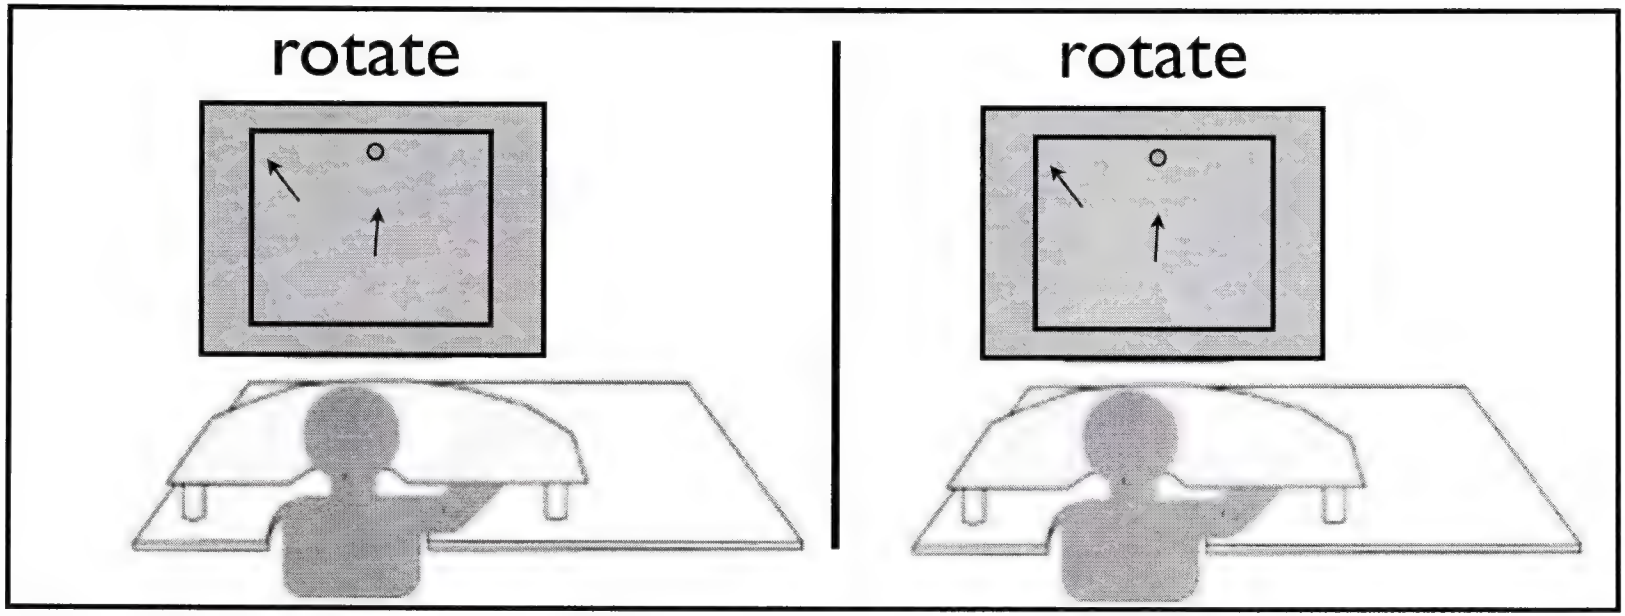

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

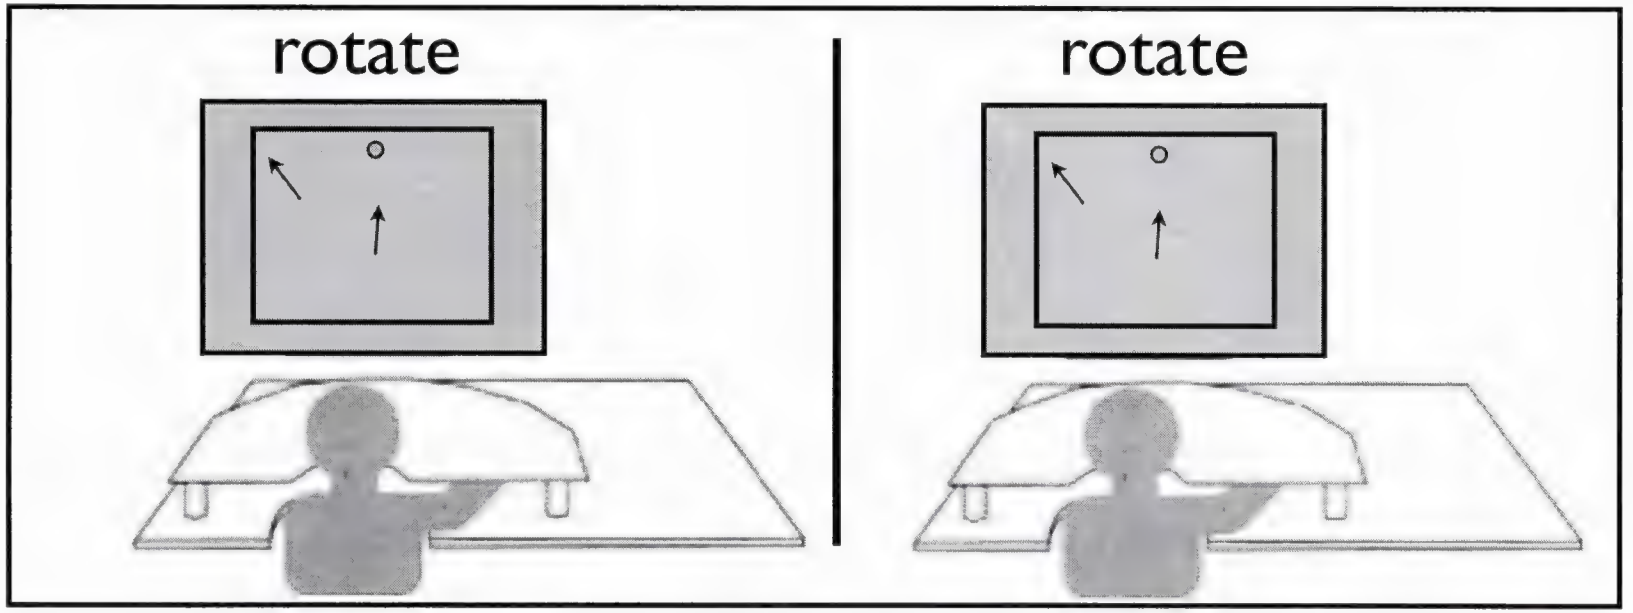

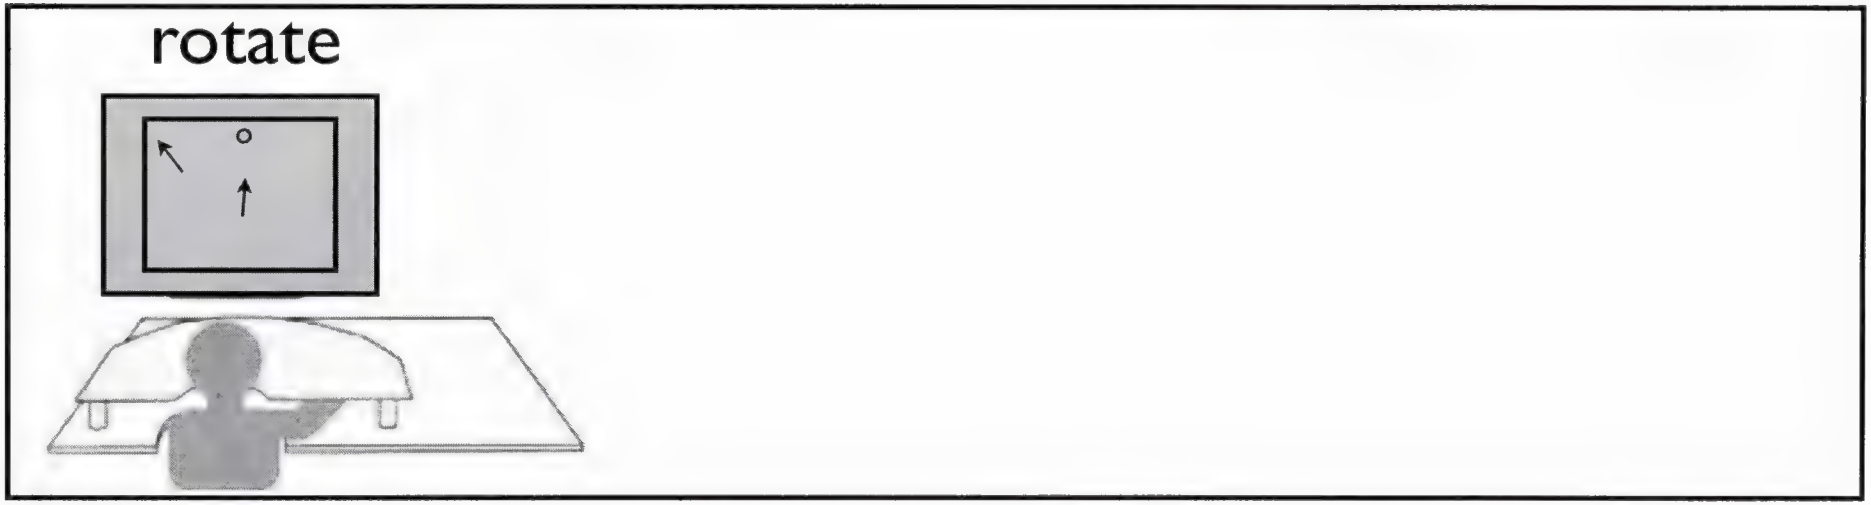

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

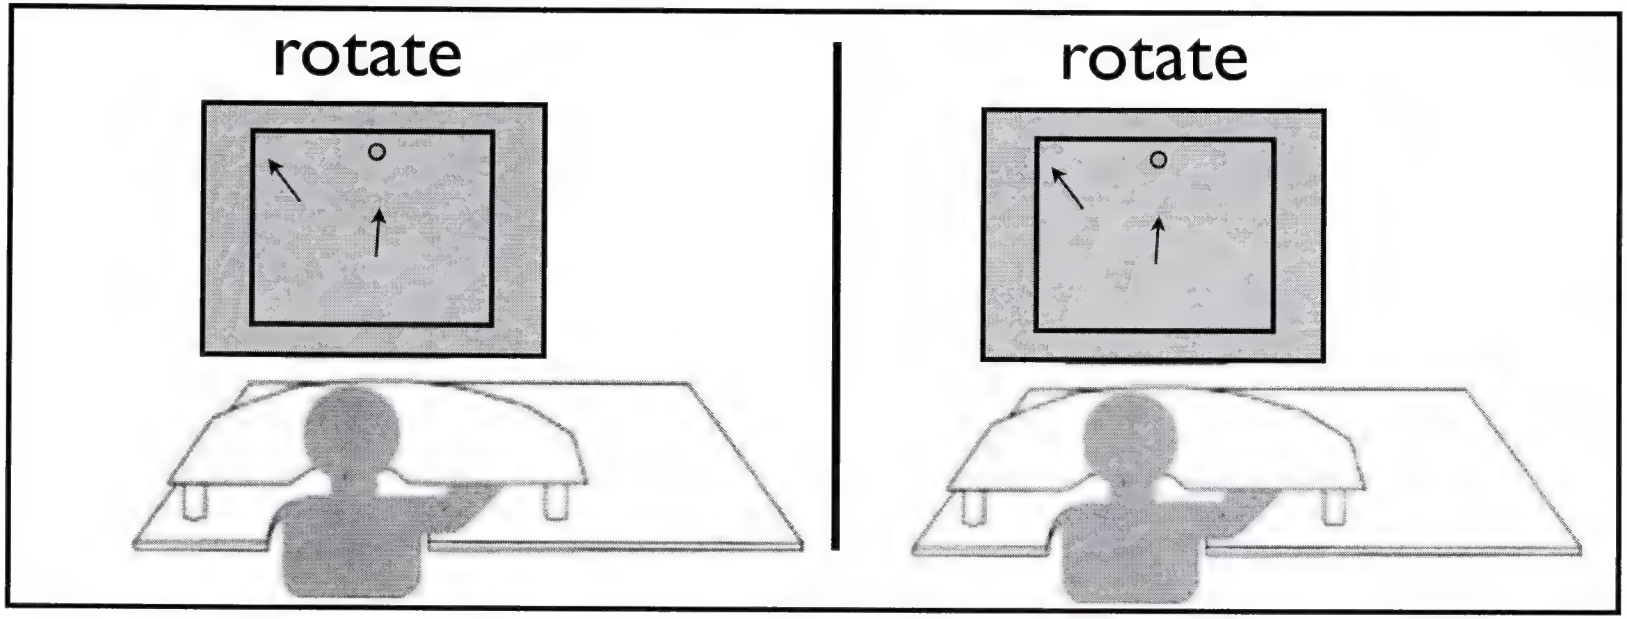

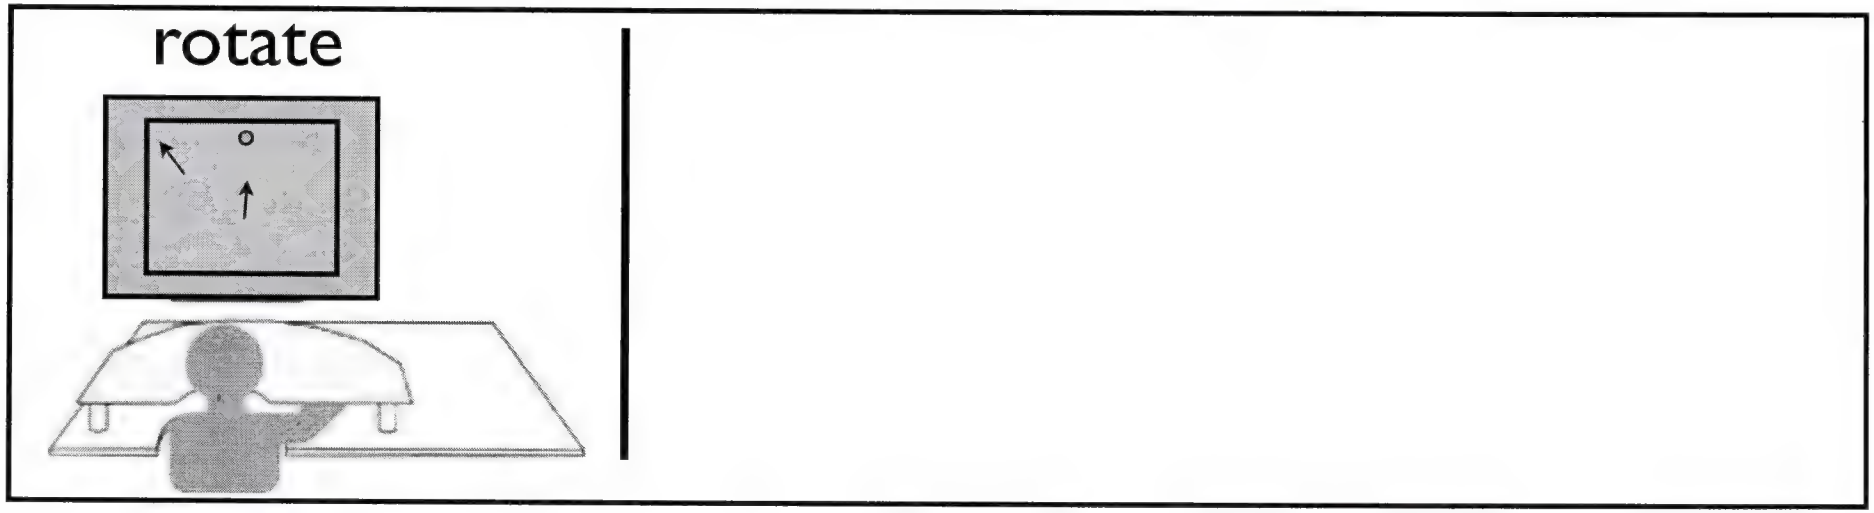

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

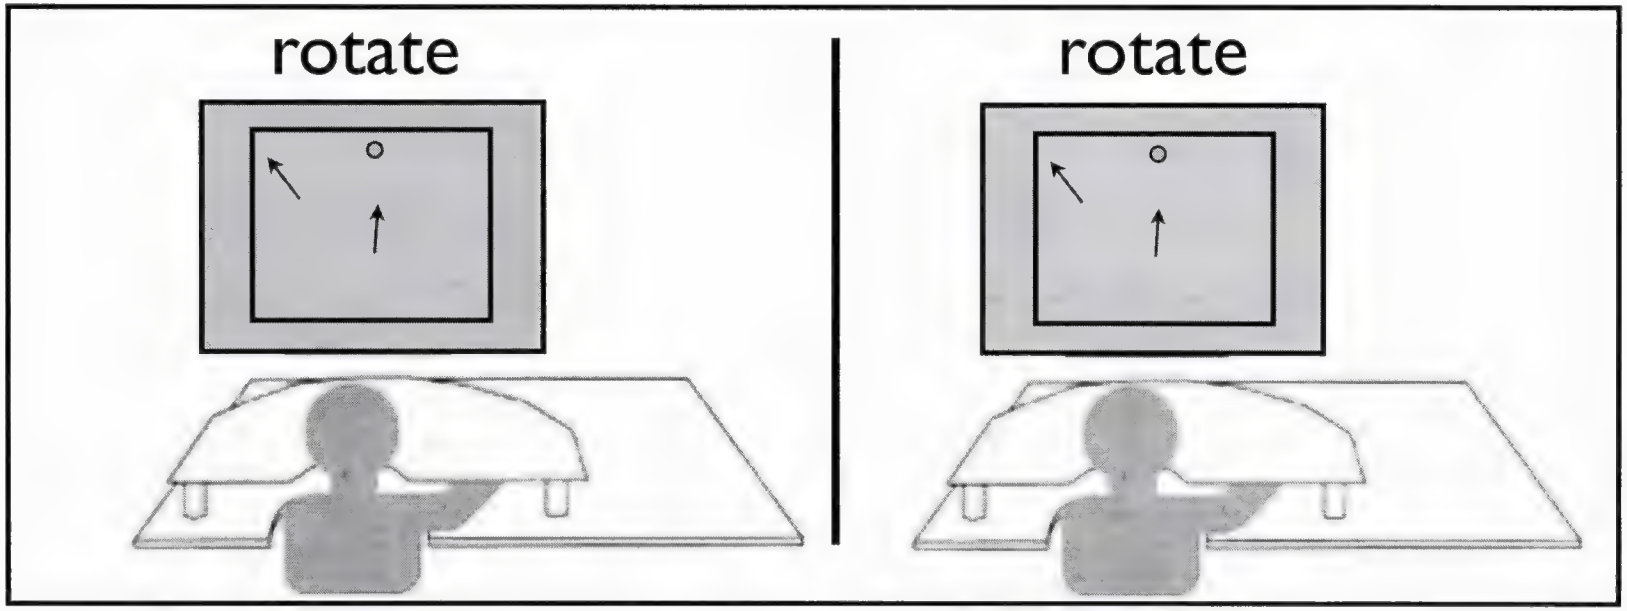

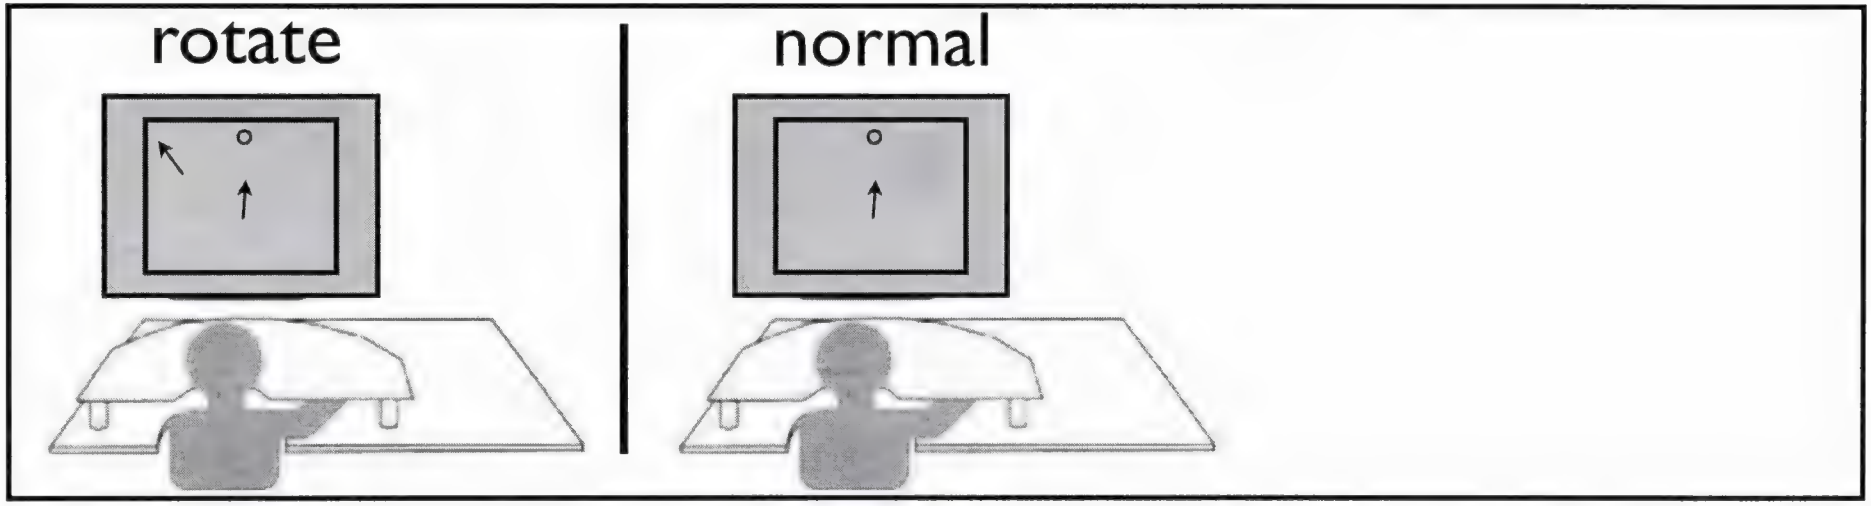

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

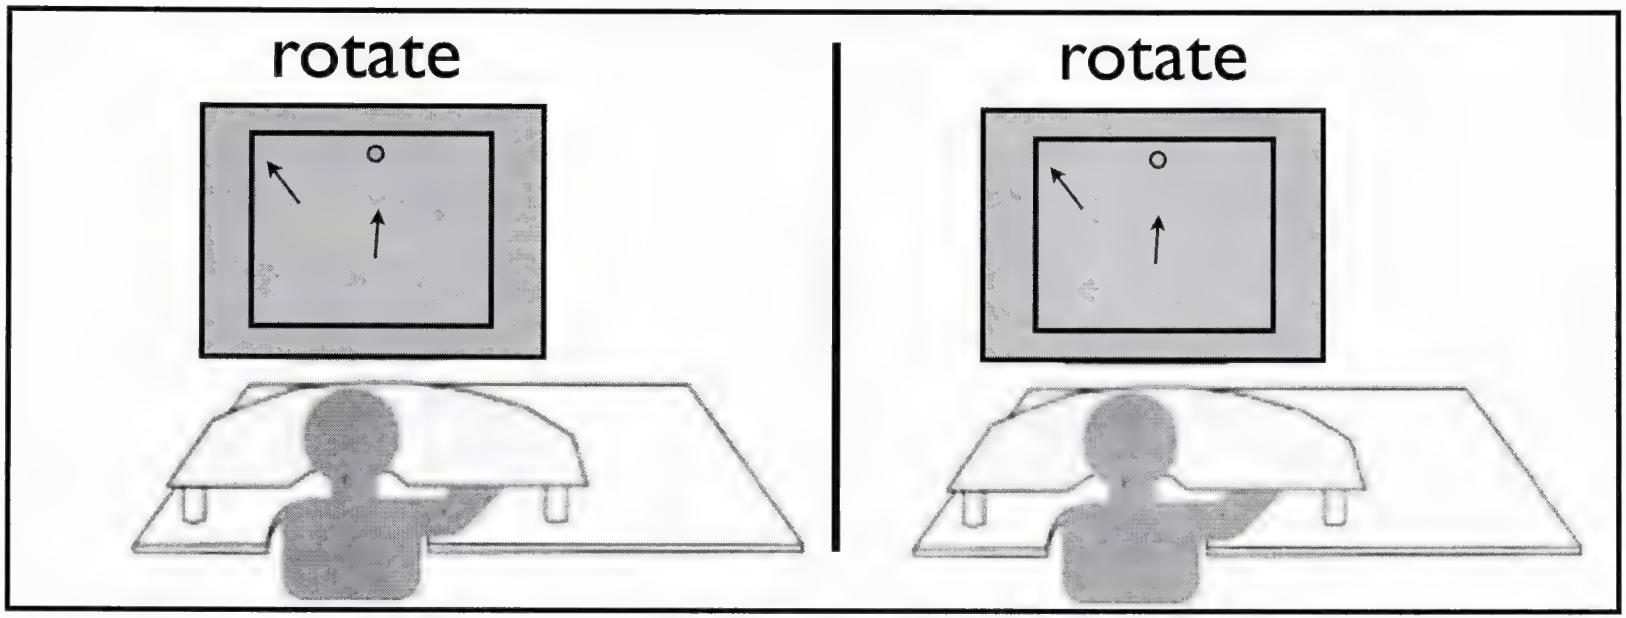

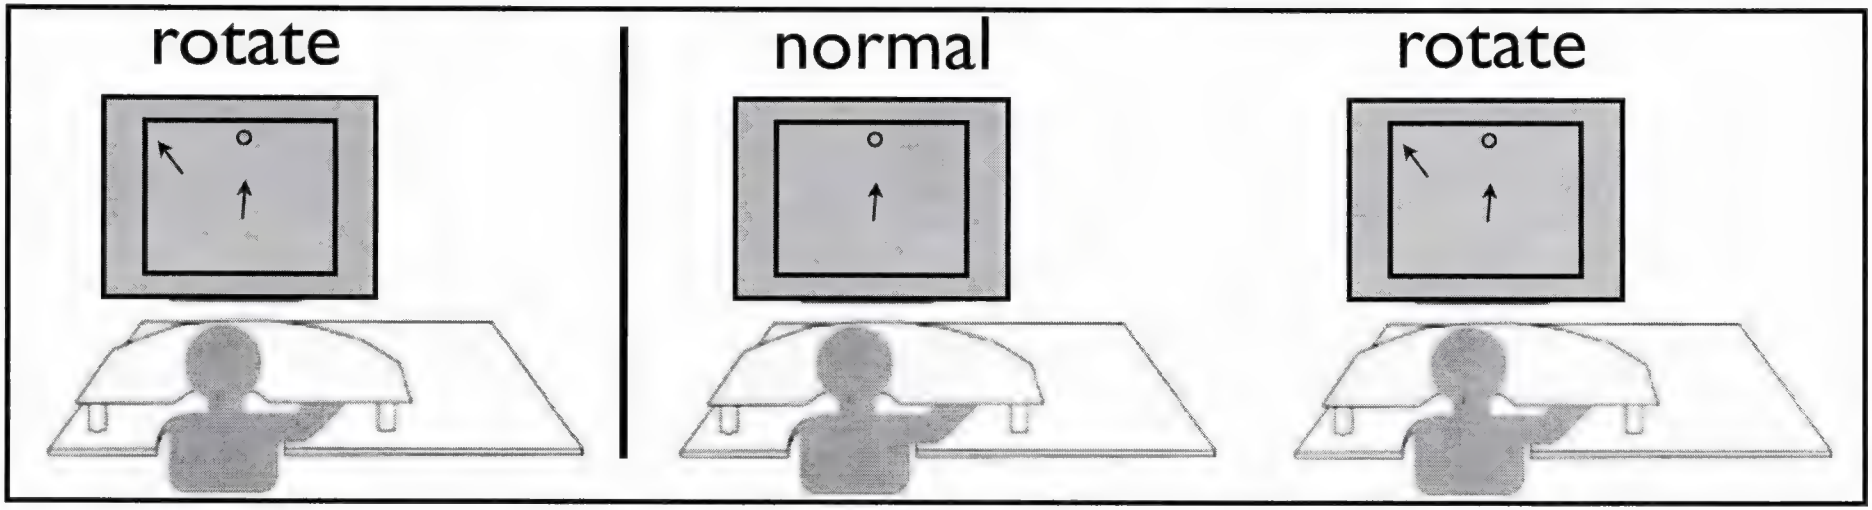

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

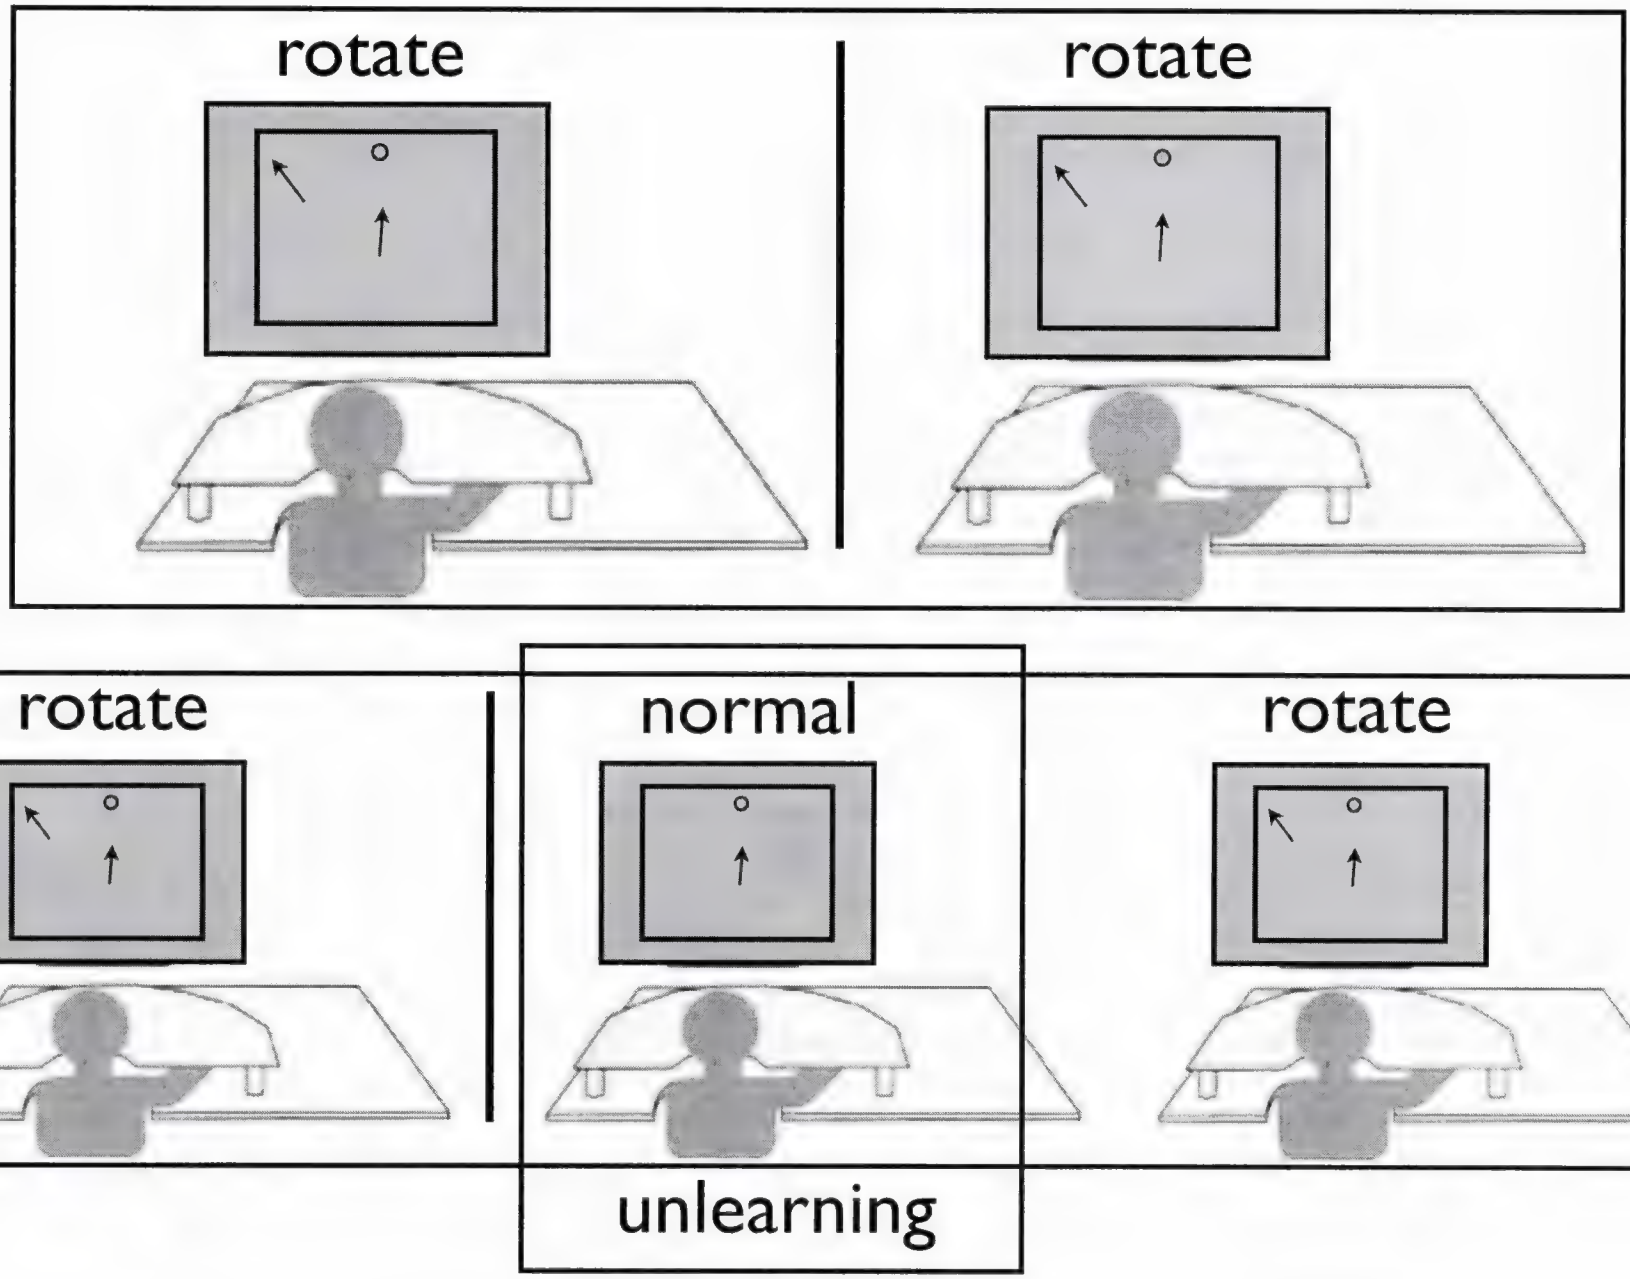

Data from Krakauer, Ghez, Ghilardi, 2005
Cartoon adapted from Wigmore et.al 2002
Smith, Ghazizadeh, Shadmehr 2006

Alex Ramirez, Stefano Fusi
John Krakauer

no unlearning

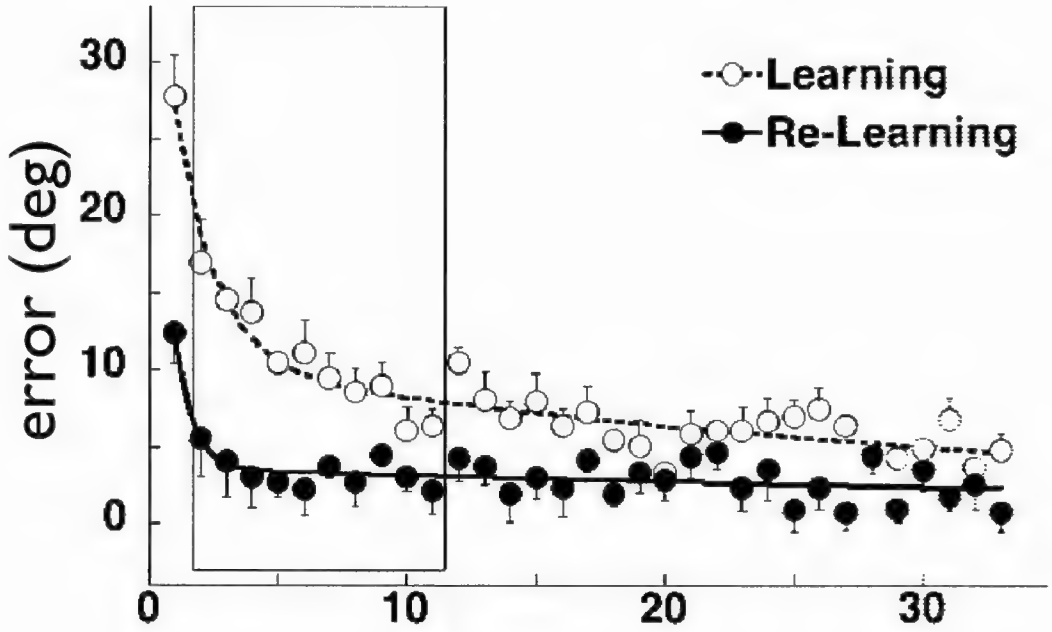

Krakauer, Ghez, Ghilardi, 2005

no unlearning

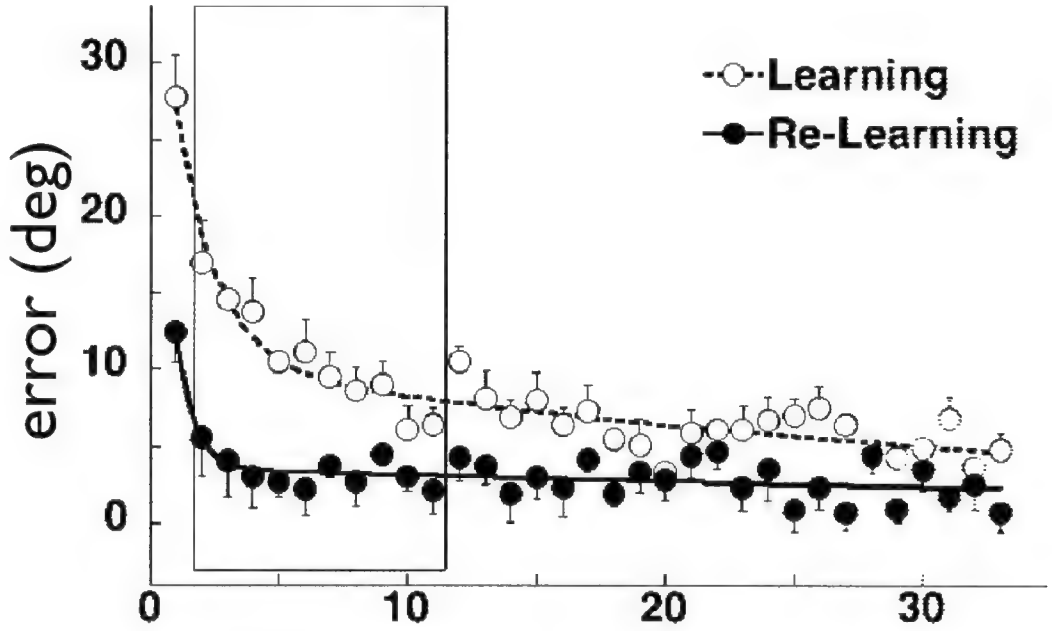

with unlearning

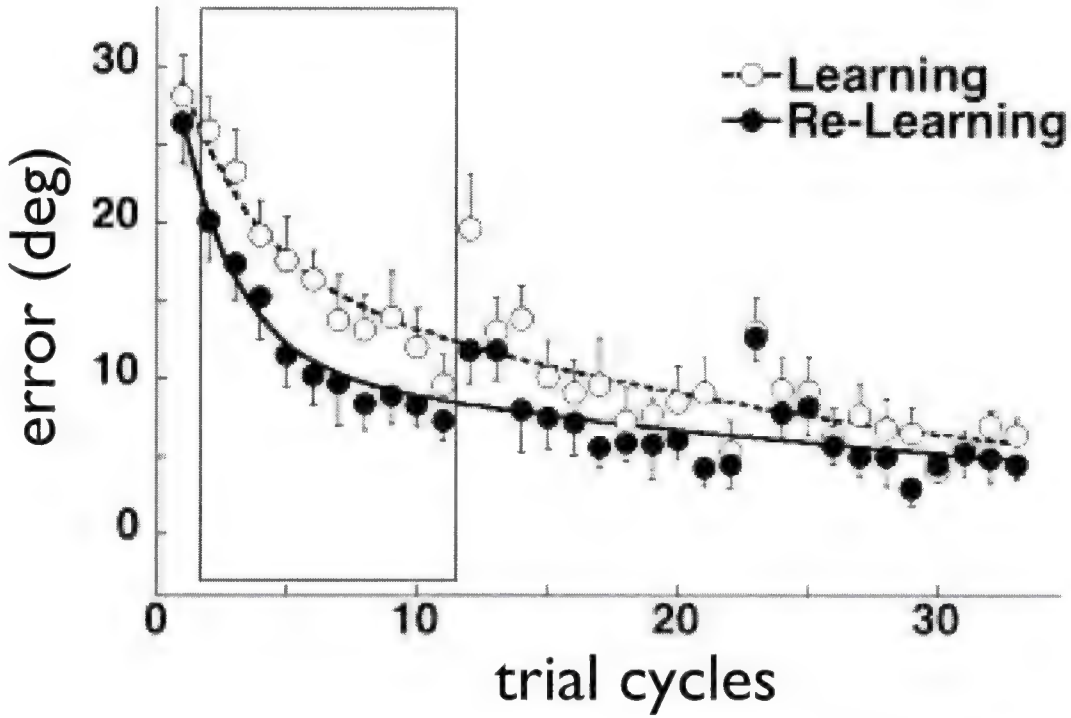

Krakauer, Ghez, Ghilardi, 2005

initial learning

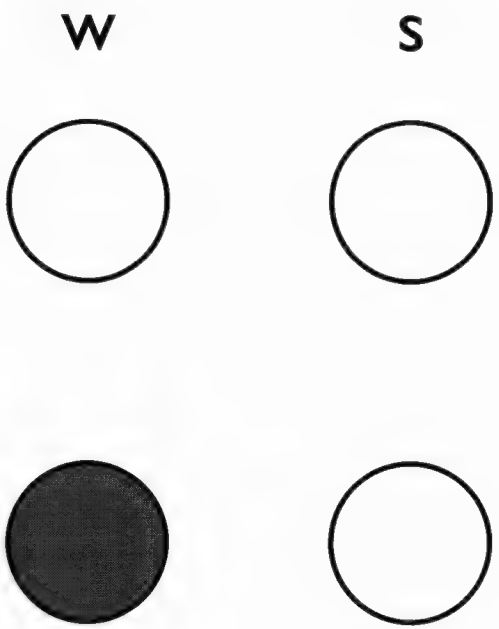

initial learning

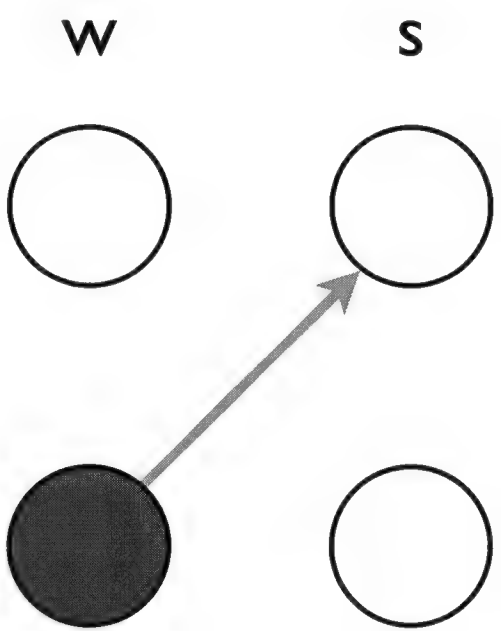

initial learning

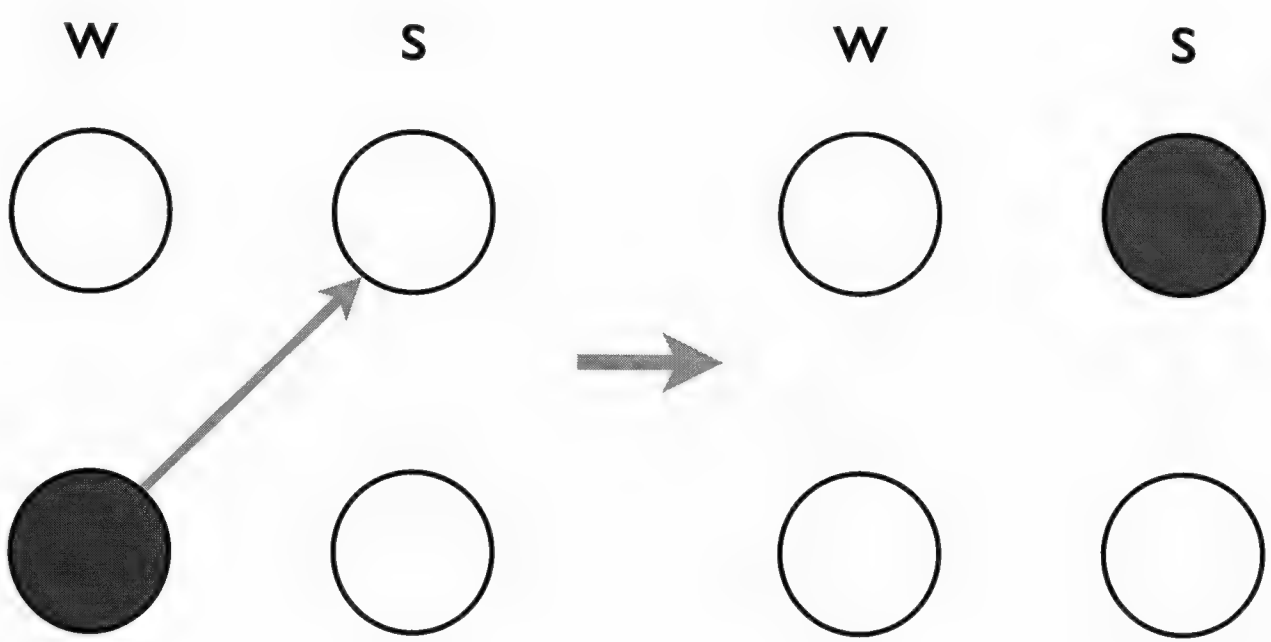

consolidation and forgetting

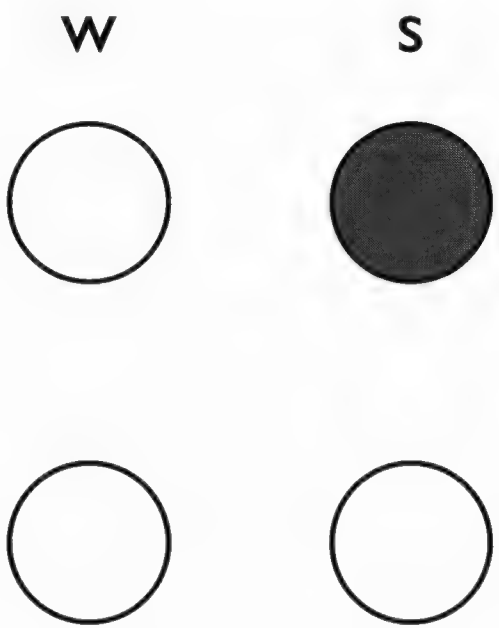

consolidation and forgetting

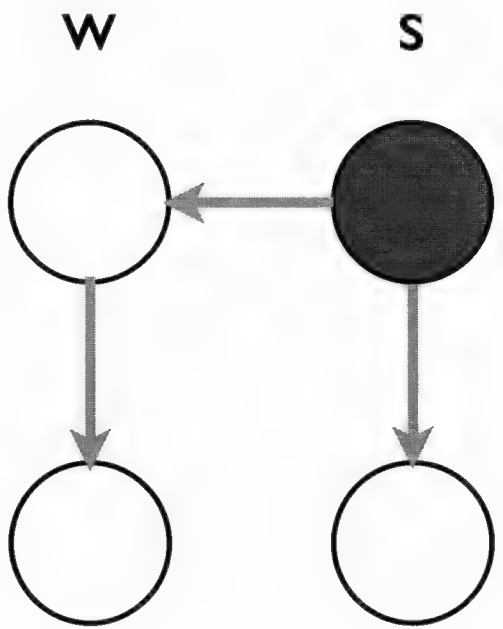

consolidation and forgetting

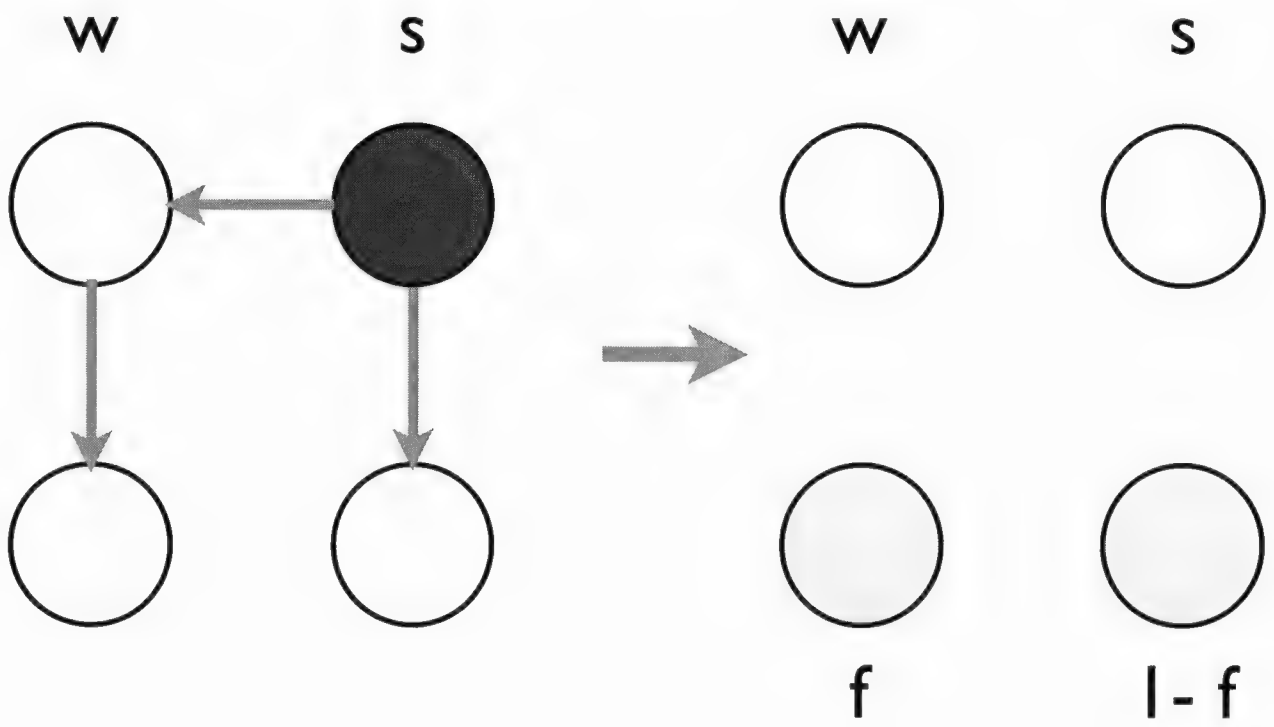

re-learning

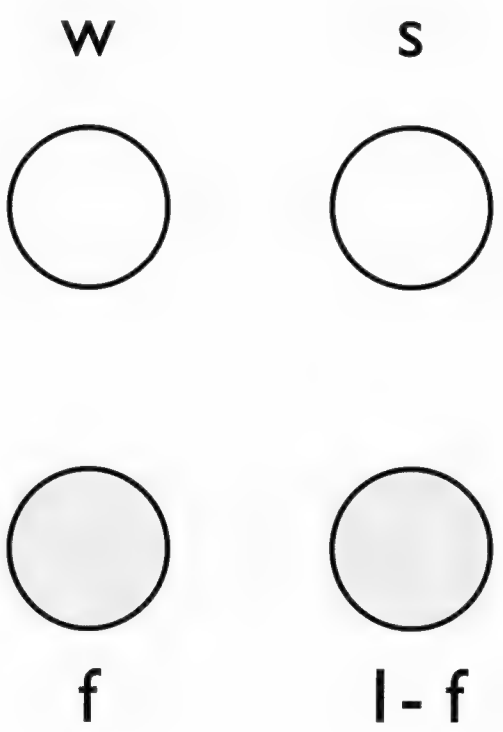

re-learning

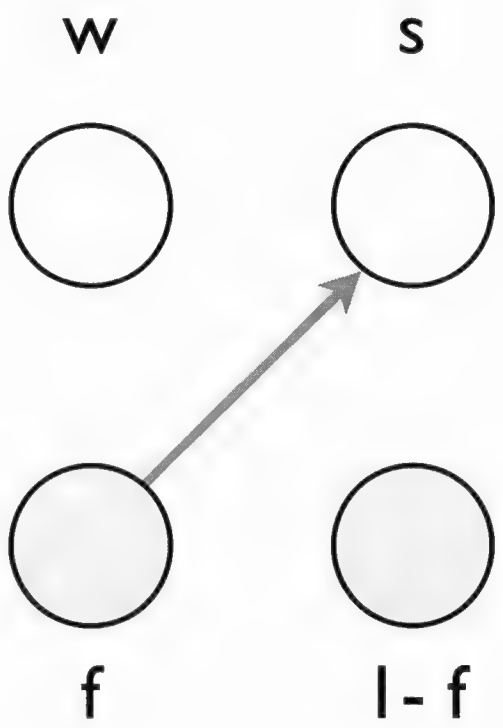

re-learning

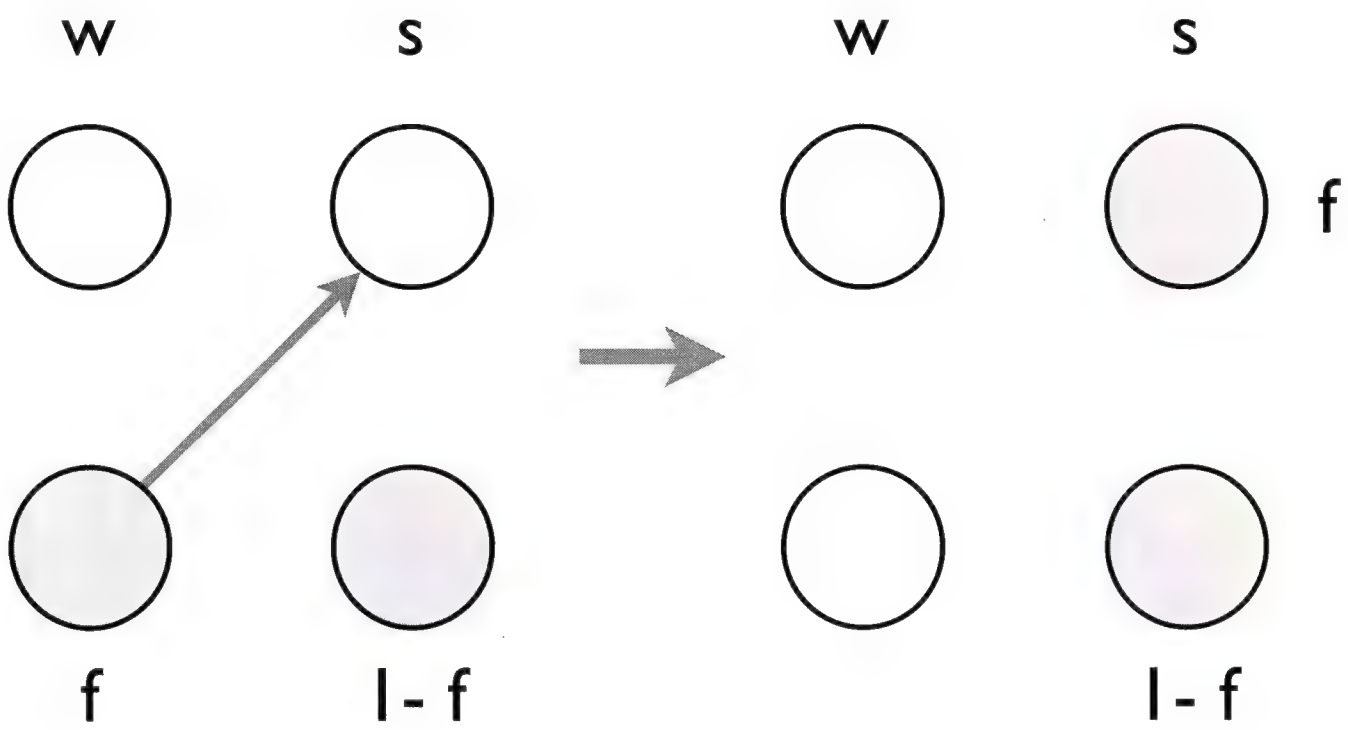

un-learning

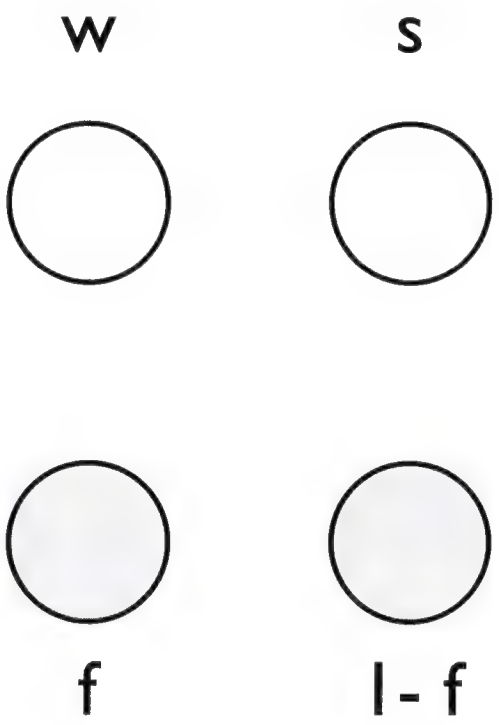

un-learning

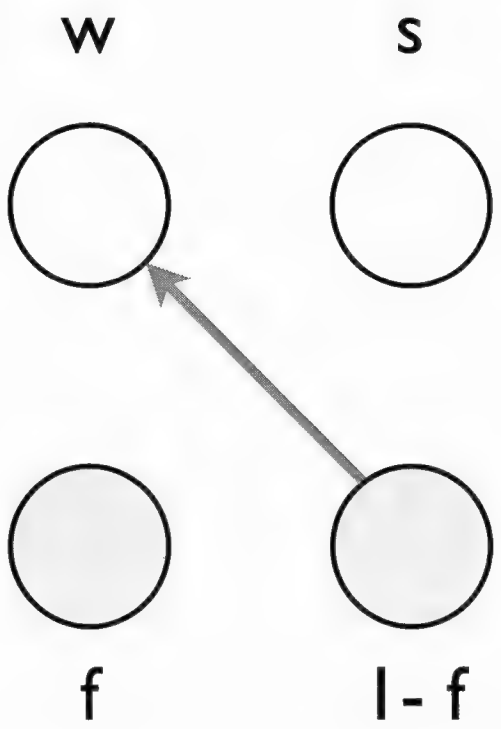

un-learning

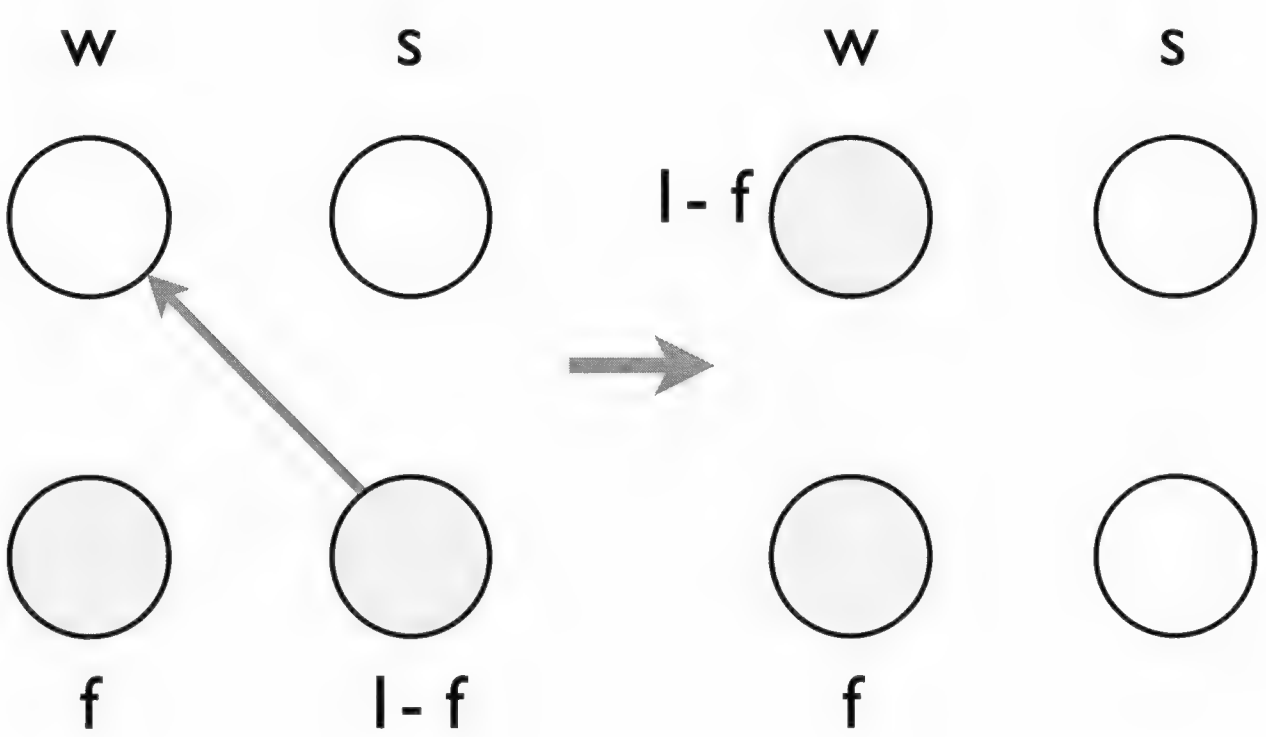

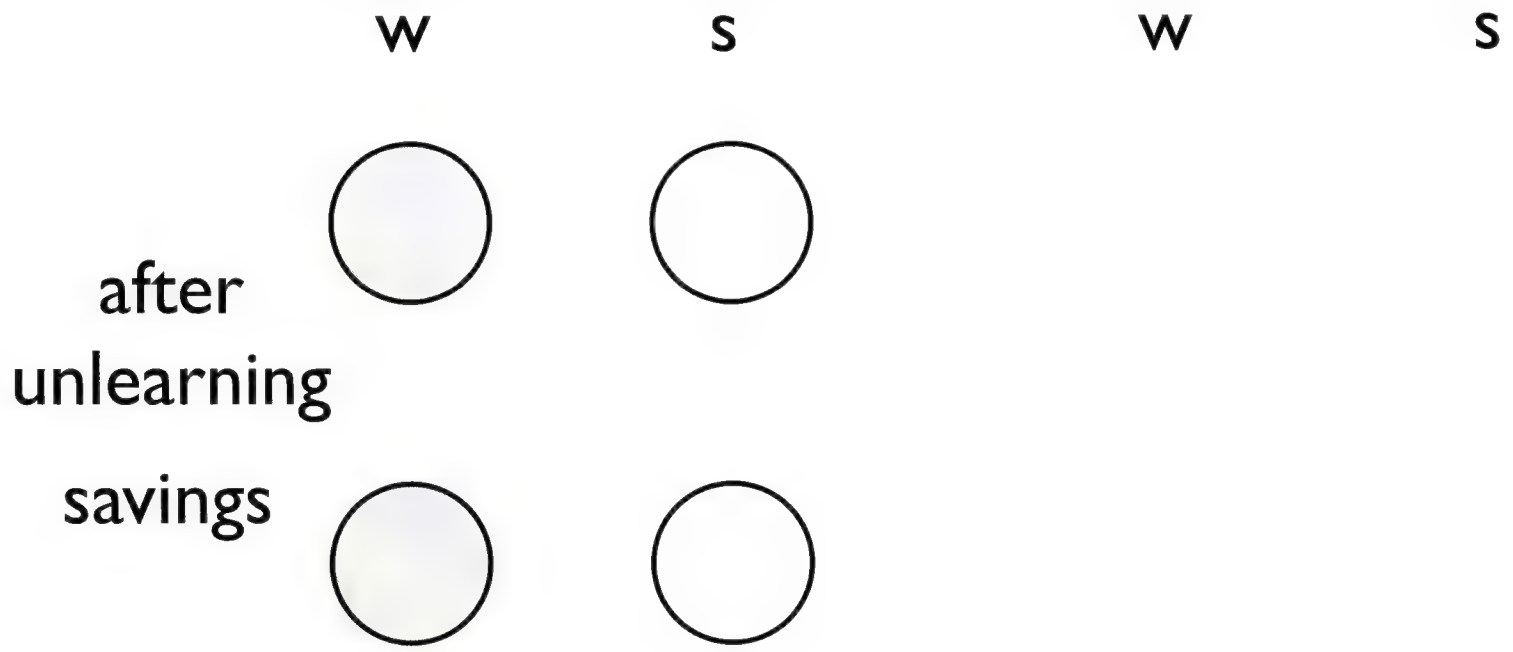
W
S
W
S
after
unlearning
savings

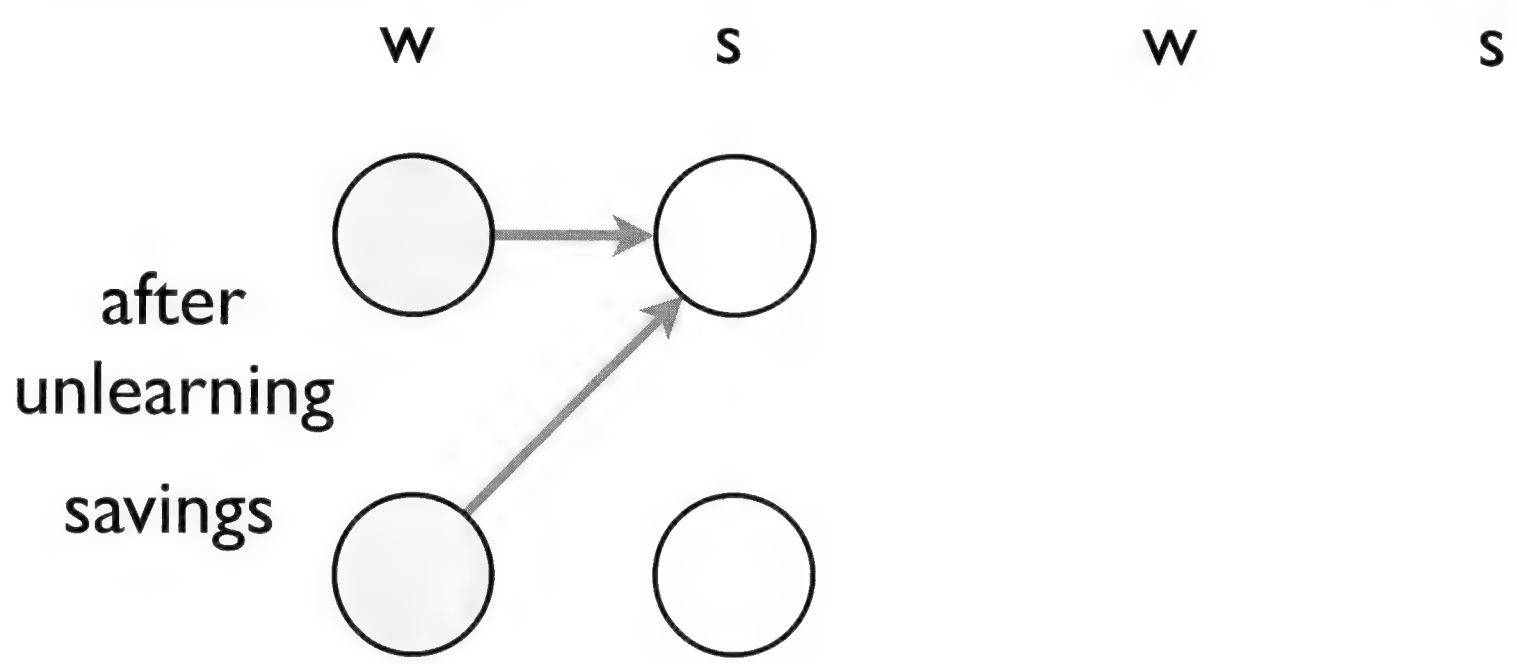
w
s
w
s
after
unlearning
savings

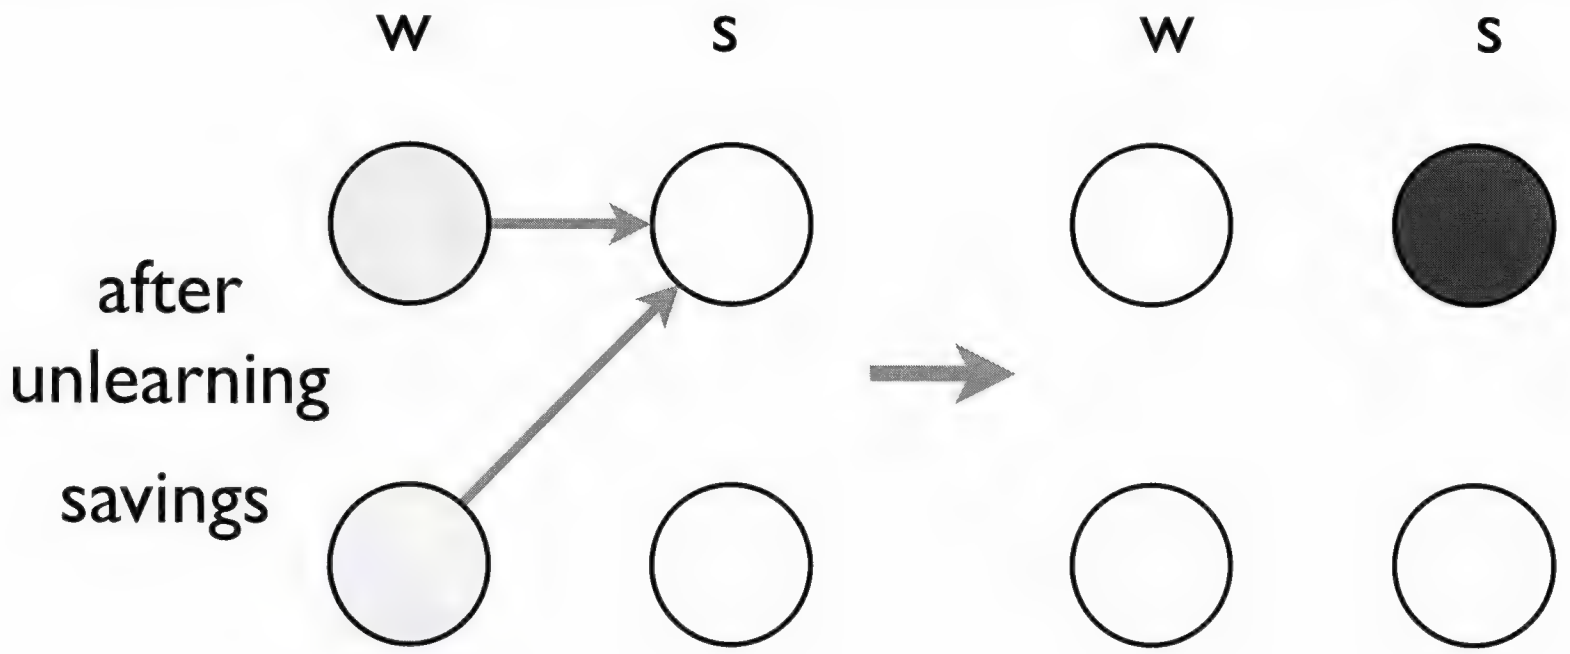
W
S
W
S
after
unlearning
savings

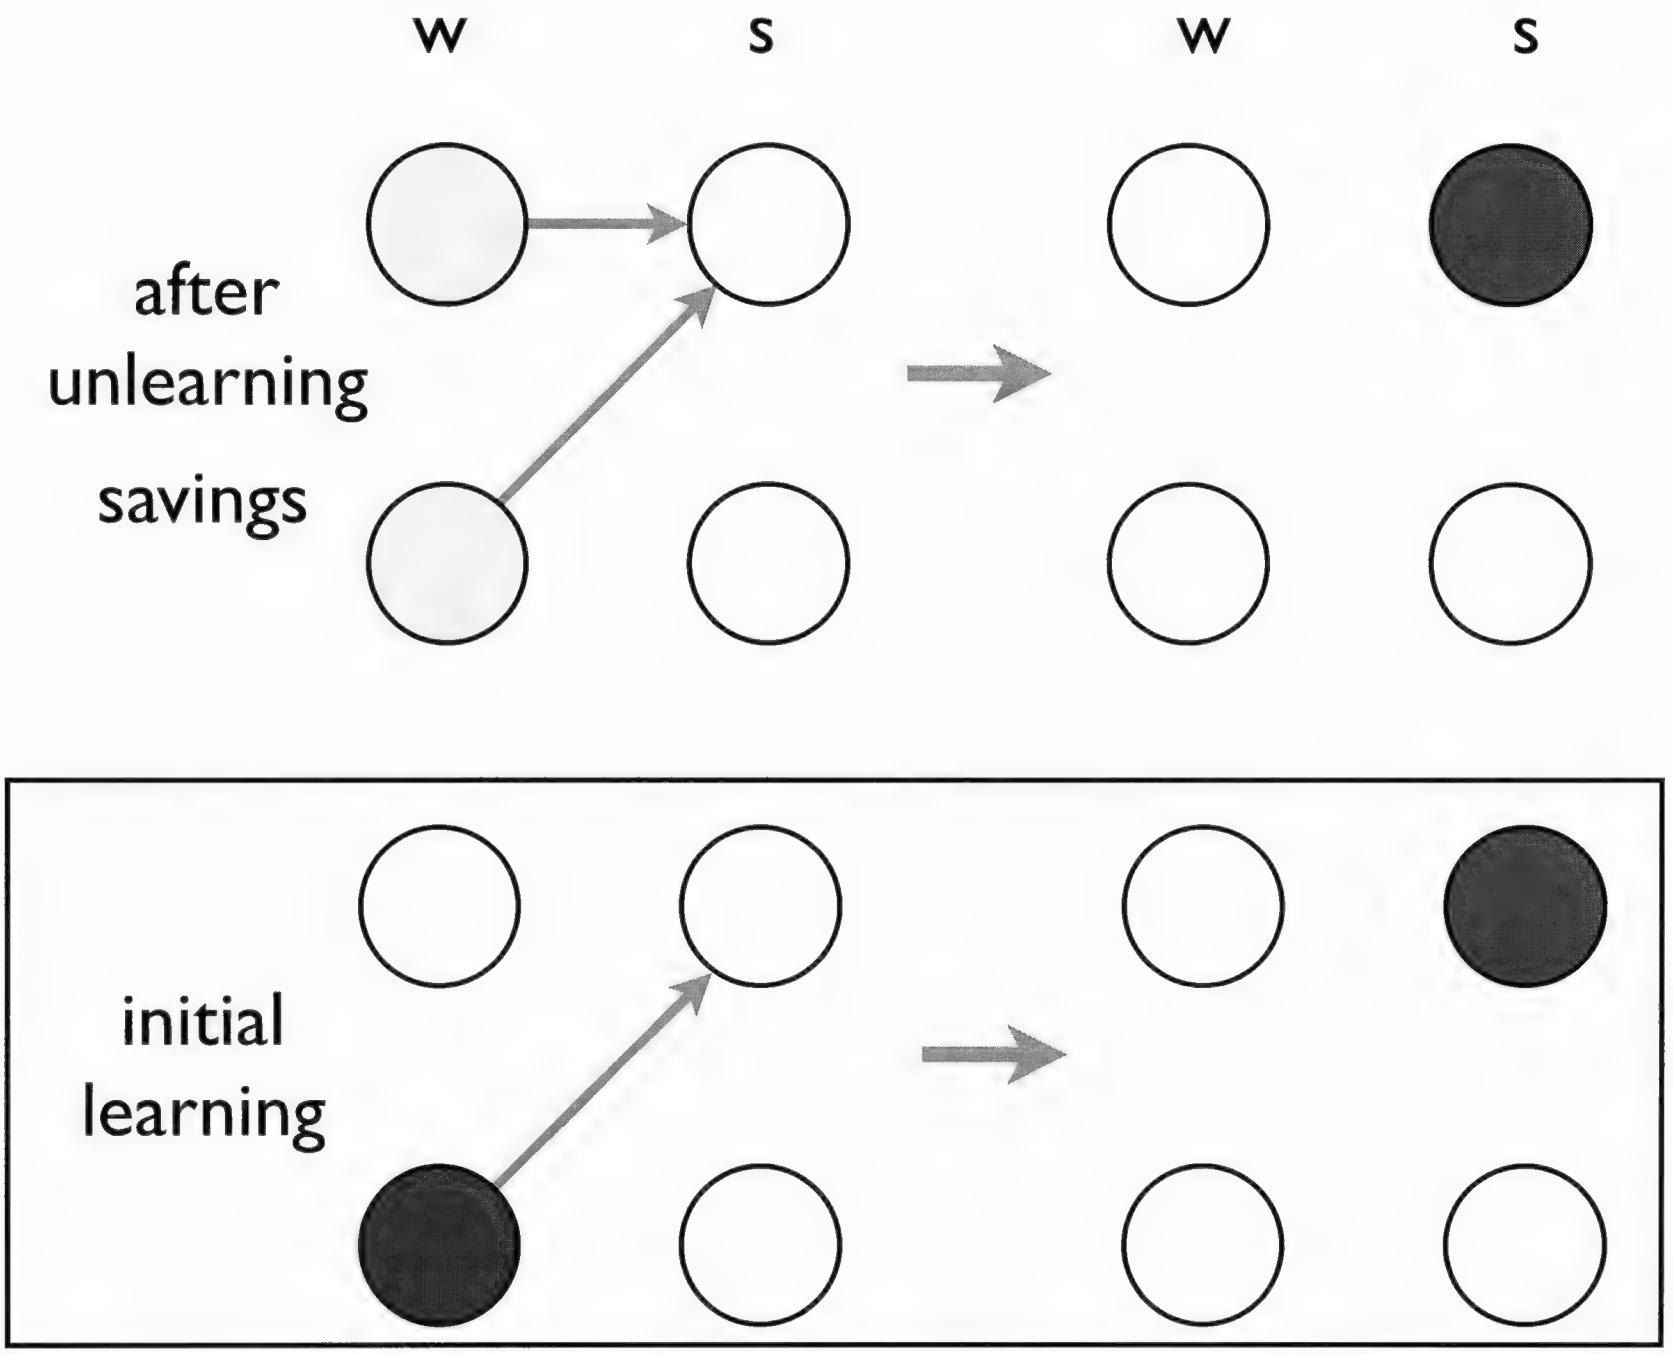
W
S
W
S
after unlearning
savings
initial learning

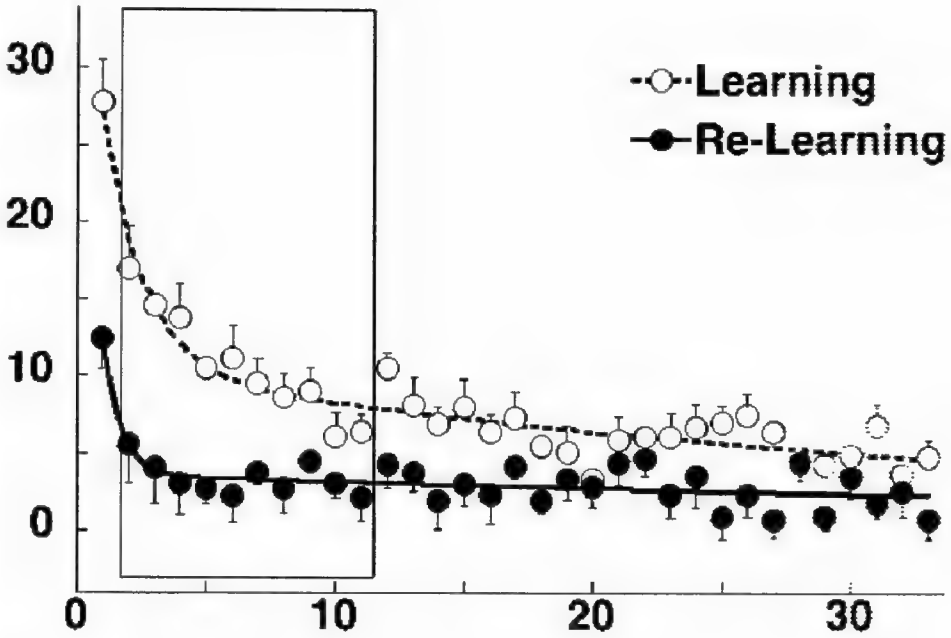

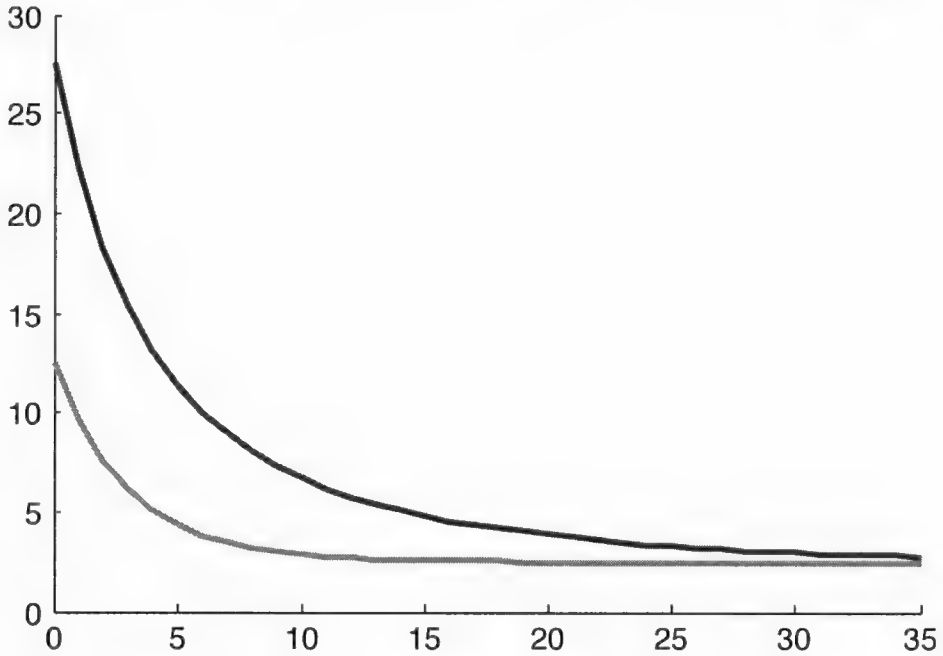

trial cycles

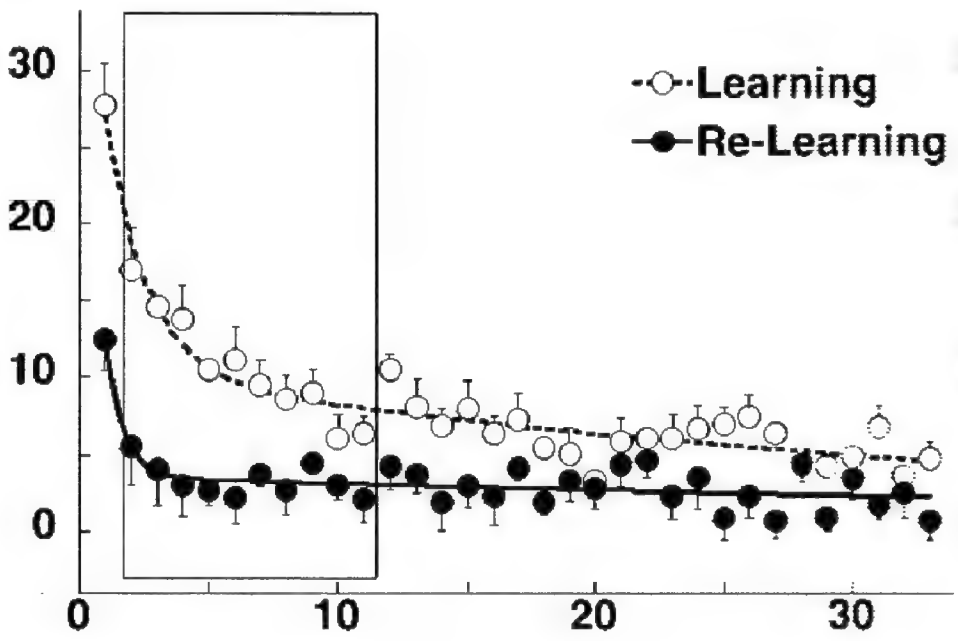
Learning
Re-Learning
30
20
10
0
0
10
20
30

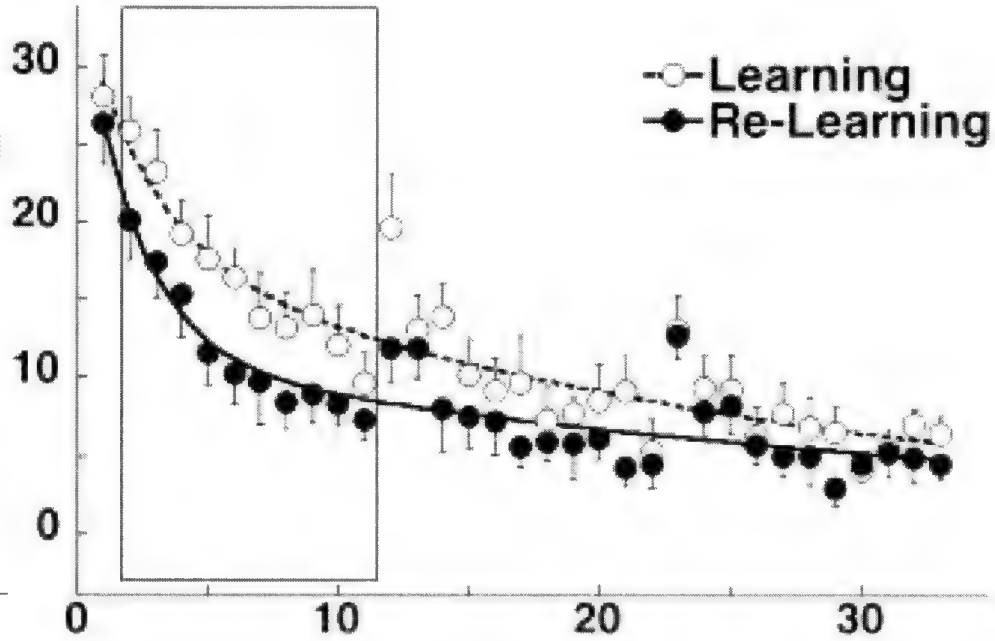
Learning
Re-Learning
30
20
10
0
0
10
20
30

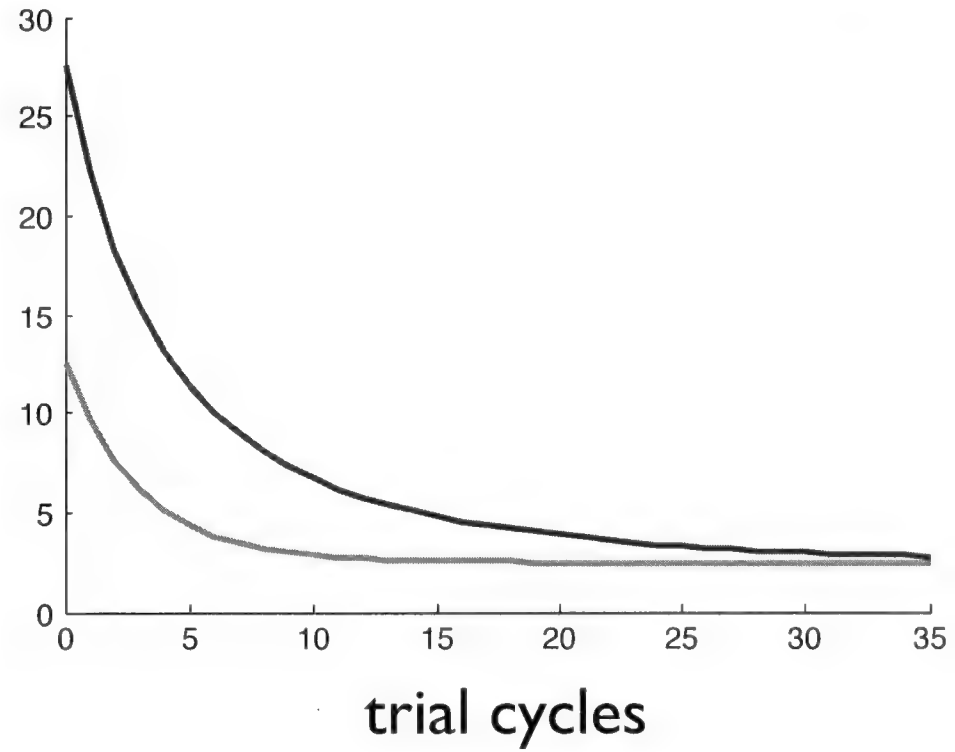
30
25
20
15
10
5
0
0
5
10
15
20
25
30
35
trial cycles

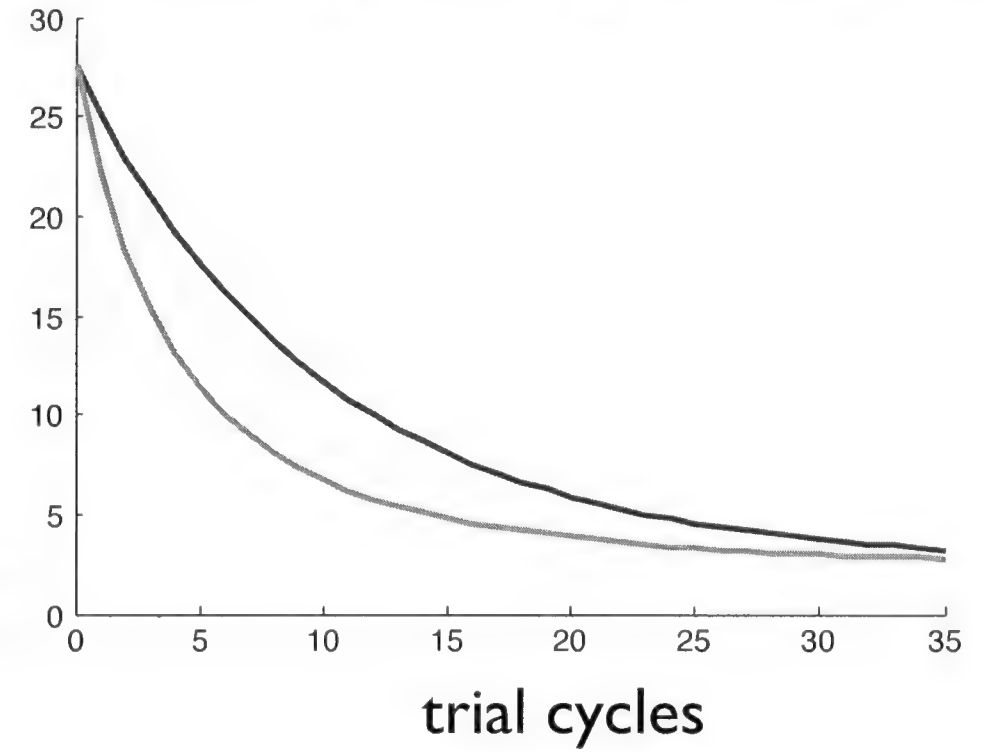
30
25
20
15
10
5
0
0
5
10
15
20
25
30
35
trial cycles

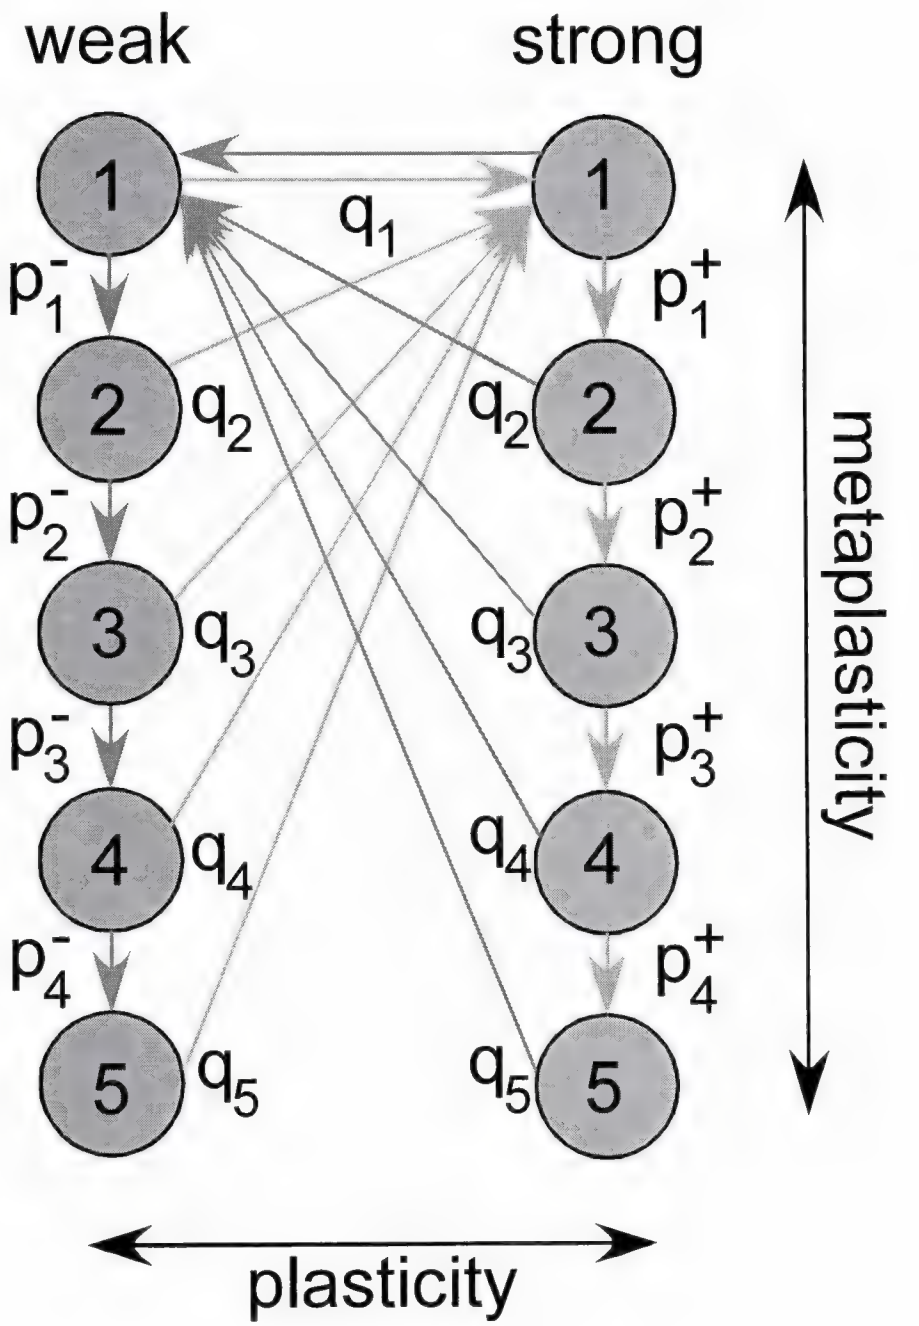
weak
strong
1
1
q_1
p_1^-
p_1^+
2
q_2
q_2
2
p_2^-
p_2^+
3
q_3
q_3
3
p_3^-
p_3^+
4
q_4
q_4
4
p_4^-
p_4^+
5
q_5
q_5
5
metaplasticity
plasticity

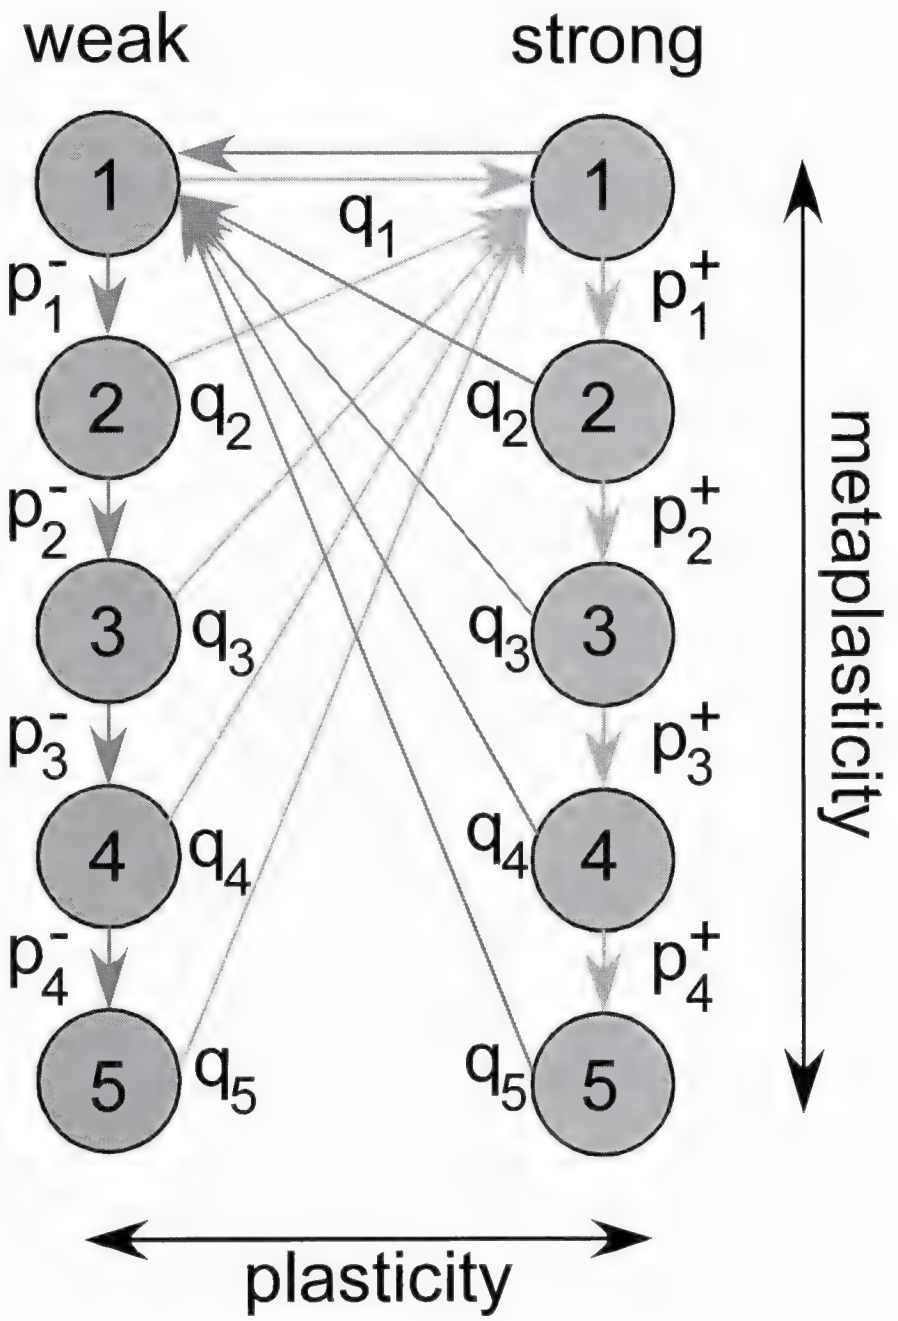

high plasticity for learning
low plasticity for retention

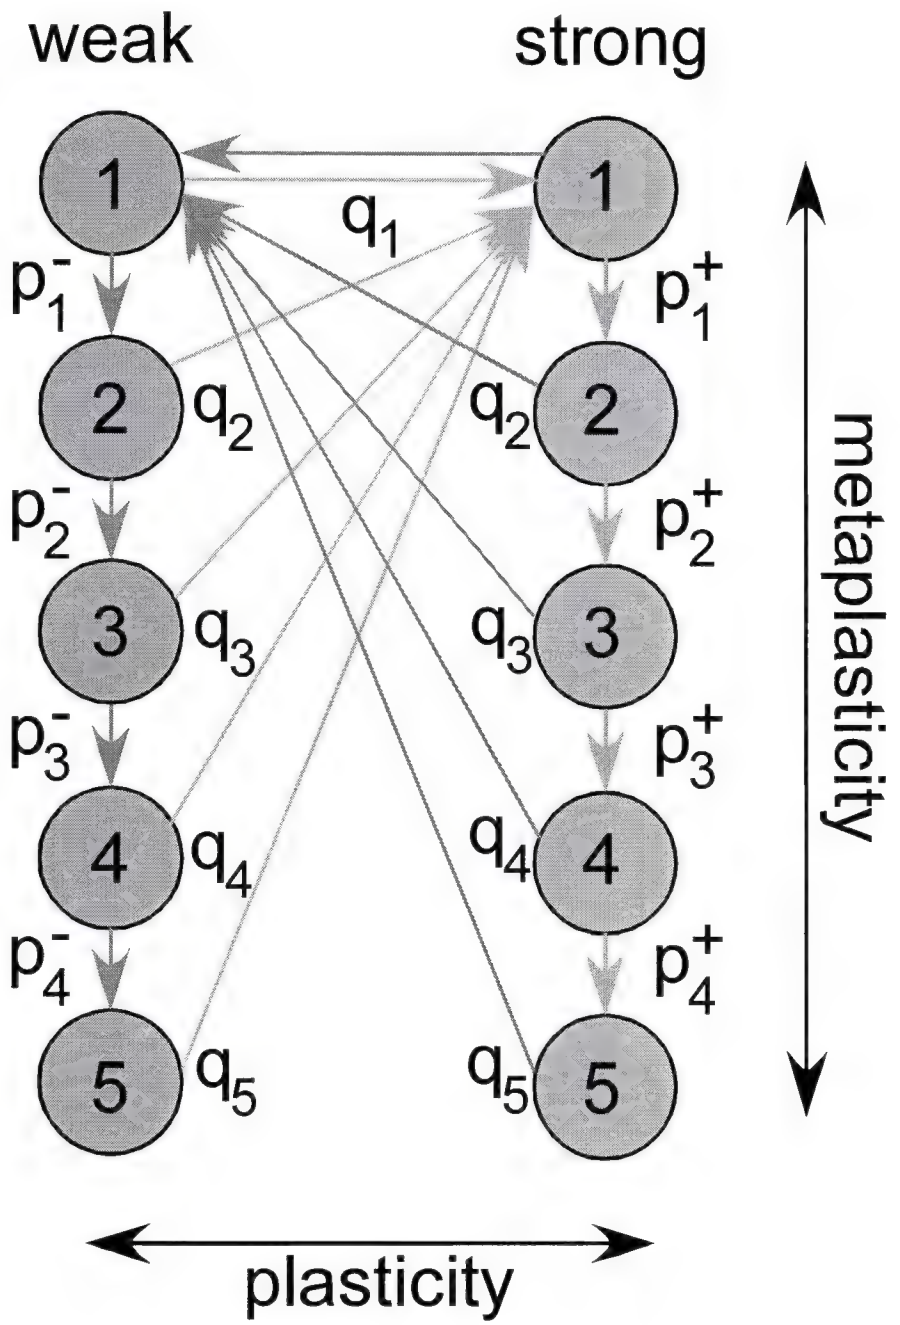

high plasticity for learning
low plasticity for retention

single-trial learning

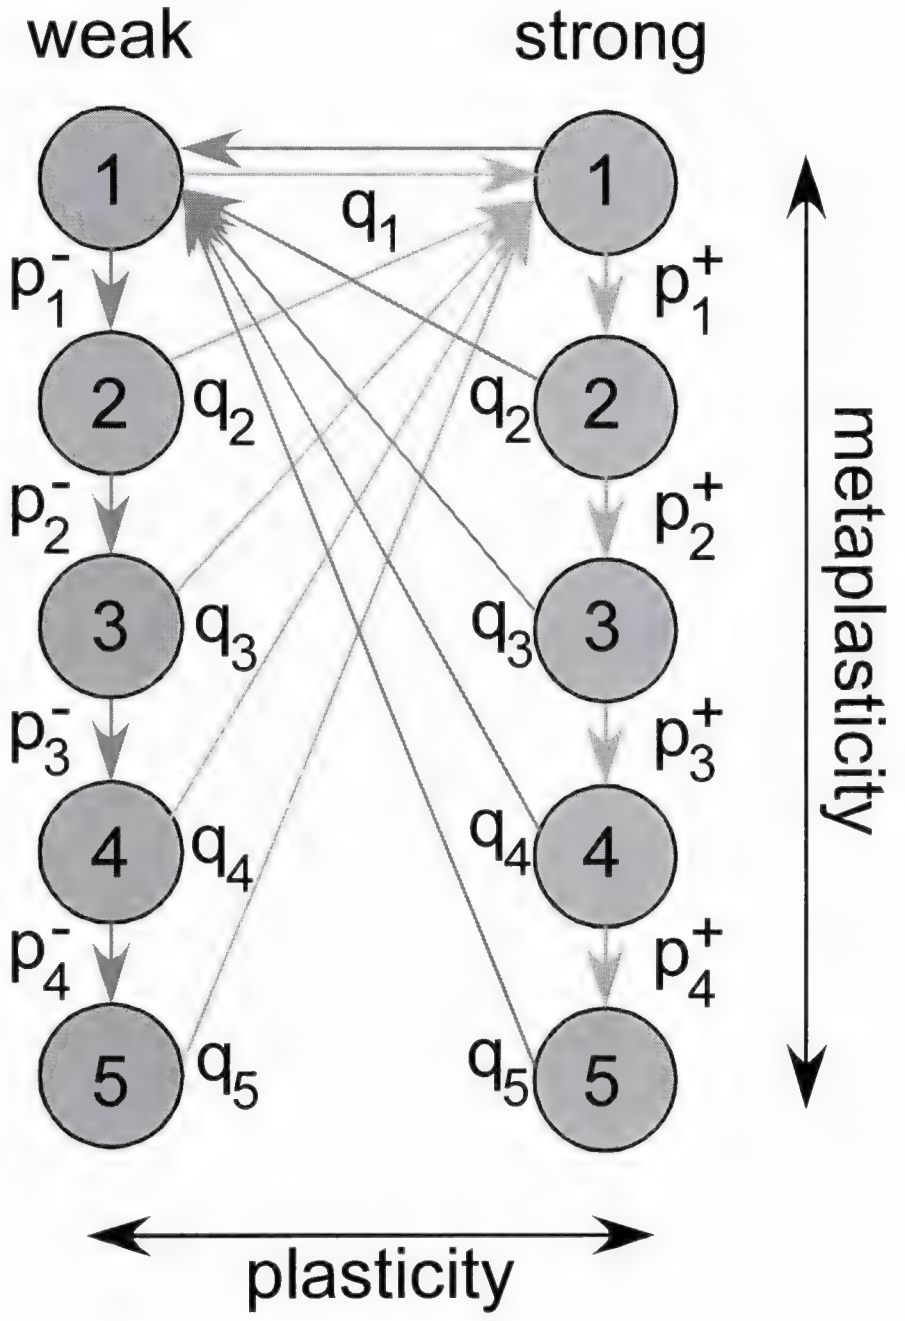

high plasticity for learning
low plasticity for retention

single-trial learning

spaced/massed effect

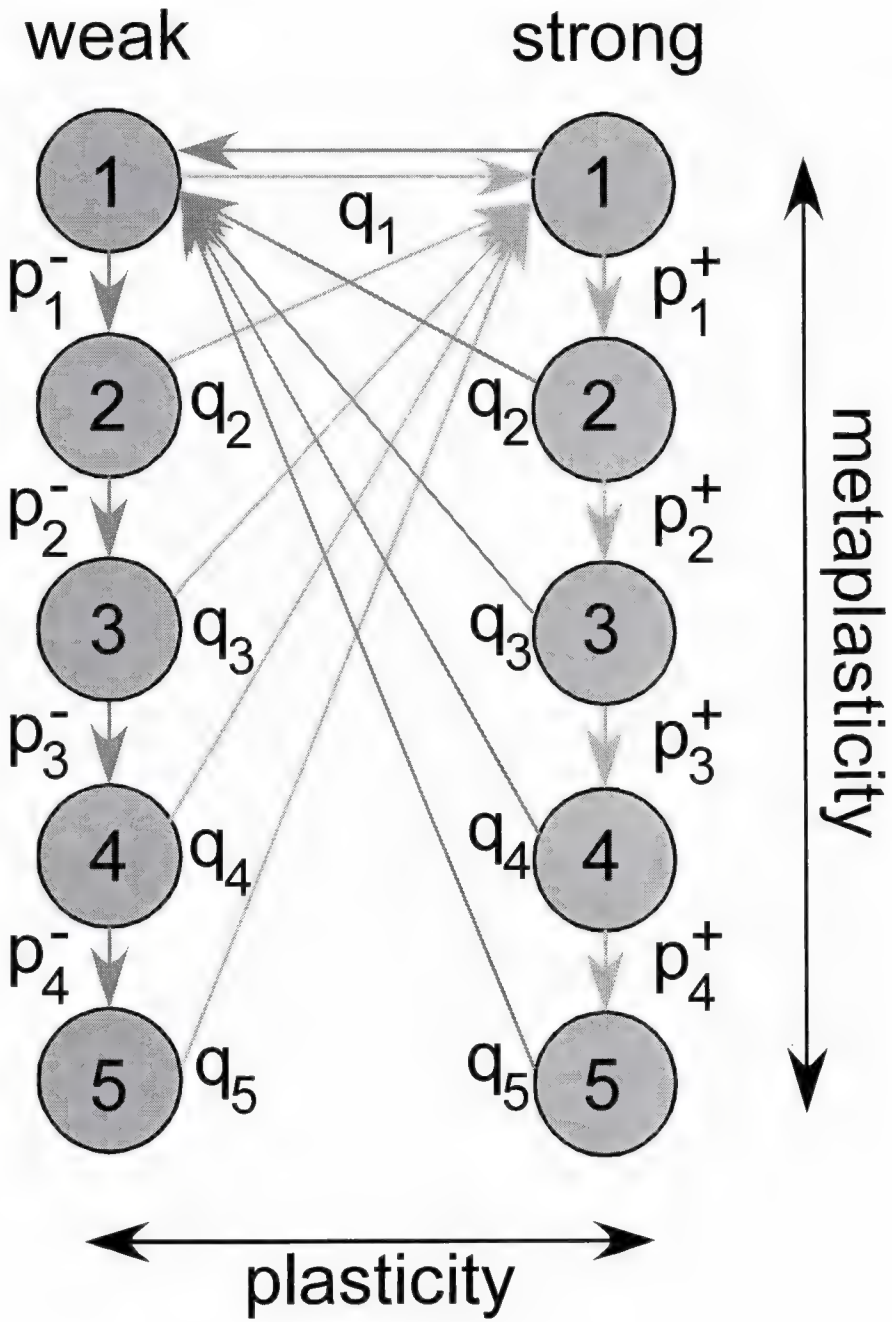

high plasticity for learning
low plasticity for retention

single-trial learning

spaced/massed effect

reinforcement learning

Generating Coherent Patterns of Activity from Chaotic Neural Networks

David Sussillo and L. F. Abbott

Department of Neuroscience
Department of Physiology and Cellular Biophysics
Columbia University College of Physicians and Surgeons
New York NY 10032-2695 USA

Abstract

Neural circuits display complex activity patterns both spontaneously and when responding to a stimulus or generating a motor output. How are these two forms of activity related? We develop a procedure called FORCE learning for modifying synaptic strengths either external to or within a model neural network to change chaotic spontaneous activity into a wide variety of desired activity patterns. FORCE learning works even though the networks we train are spontaneously chaotic and we leave feedback loops intact and unclamped during learning. Using this approach, we construct networks that produce a wide variety of complex output patterns, input-output transformations that require memory, multiple outputs that can be switched by control inputs, and motor patterns matching human motion capture data. Our results reproduce data on pre-movement activity in motor and premotor cortex, and suggest that synaptic plasticity may be a more rapid and powerful modulator of network activity than generally appreciated.

Introduction

When we voluntarily move a limb or perform some other motor action, what is the source of the neural activity that initiates and carries out this behavior? We explore the idea that such actions arise from the reorganization of spontaneous neural activity. This hypothesis raises another question: How can apparently chaotic spontaneous activity be reorganized into the coherent patterns required to generate controlled actions? Following, but modifying and extending, earlier work (Jaeger and Haas, 2004; Maass et al., 2007), we show how external feedback loops or internal synaptic modifications can be used to alter the chaotic activity of a recurrently connected neural network and generate complex but controlled outputs.

Training a neural network is a process through which network parameters (typically synaptic strengths) are modified on the basis of output errors until a desired response is produced. Researchers in the machine learning and computer vision communities have developed powerful methods for training artificial neural networks to perform complex tasks (Rumelhart and McClelland, 1986; Hinton et al., 2007), but these apply predominantly to networks with feedforward architectures. Biological networks tend to be connected in a highly recurrent manner. Training procedures have also been developed for recurrently connected neural networks (Rumelhart et al., 1986; Williams and Zipser, 1989; Pearlmutter, 1989; Atiya and Parlos, 2000), but these are more computationally demanding and difficult to use than feedforward learning algorithms, and there are fundamental limitations to their applicability (Doya, 1992). In particular, these algorithms generally will not converge if applied to recurrent neural networks with chaotic activity, that is, activity that is irregular and exponentially sensitive to initial conditions (Abarbanel et al., 2008). This limitation is severe because models of spontaneously active neural circuits typically exhibit chaotic dynamics. For example, spiking models of spontaneous activity in cortical circuits (van Vreeswijk and Sompolinsky, 1996; Amit and Brunel, 1997; Brunel, 2000), which can generate realistic patterns of activity, and the analogous spontaneously active firing-rate model networks that we use here have chaotic dynamics (Sompolinsky et al., 1988).

To develop a successful training procedure for recurrent neural networks, we must solve three problems. First, feeding erroneous output back into a network during training can cause its activity to deviate so far from what is needed that learning fails to converge. In previous work (Jaeger and Haas, 2004), this problem was avoided by removing all errors from the signal fed back into the network. In addition to the usual synaptic modification, this scheme required a mechanism for removing feedback errors, and it is difficult to see how this latter requirement could be met in a biological system. Furthermore, feeding back a signal that is identical to the desired network output prevents the network from sampling fluctuations during training, which can lead to stability problems in the final network. Here, we show how the synaptic modification procedure itself can be used to control the feedback signal, without any other mechanism being required, in a manner that allows fluctuations to be sampled and stabilized. For reasons given below, we call this procedure FORCE learning.

A second problem with training that is particularly severe in recurrent networks is credit assignment for output errors. Credit assignment amounts to figuring out which neurons and synapses are most responsible for output errors and therefore most in need of modification. This problem is particularly challenging for network units that do not produce the output directly. Jaeger and Haas (2004) dealt with this issue by restricting modification solely to synapses directly driving network output. Initially, we follow their lead in this, using the architecture of figure 1A. However, the power of FORCE learning allows us to train networks with the architectures shown in figures 1B and 1C, in which modifications are not restricted to network outputs. For reasons discussed below, these architectures are more biologically plausible than the network in figure 1A.

The third problem we address is training in the face of chaotic spontaneous activity. Jaeger and Haas (2004) avoided this problem by starting with networks that were inactive in the absence of input (which is the basis for calling them echo-state networks). As we show in the Results, there are significant advantages in using a network that exhibits chaotic activity prior to training. To exploit these advantages, however, we must avoid chaotic network activity during training. It turns out that the solution for learning in a recurrent network and for suppressing chaos turn out to be one and the same: synaptic modifications must be strong and rapid during the initial phases of training. This is precisely what the FORCE procedure achieves.

FORCE learning operates quite differently from traditional training in neural networks. Usually, training consists of performing a sequence of modifications that slowly reduce initially large errors in network output. In FORCE learning, errors are always small, even from the beginning of the training process. As a result, the goal of training is not significant error reduction, but rather reducing the amount of modification needed to keep the errors small. By the end of the training period, modification is no longer needed, and the network can generate the desired output autonomously.

From a machine learning point of view, the FORCE procedure we propose provides a powerful algorithm for constructing recurrent neural networks that generate complex and controllable patterns of activity either in the absence of or in response to input. From a biological perspective, it can be viewed either as a model for training-induced modification or, more conservatively, as a method for building functioning circuit models for further study. Either way, our approach introduces a novel way to think about learning in neural networks and to make contact with experimental data.

Results

The recurrent network that forms the basis of our studies is a conventional model in which the outputs of individual neurons are characterized by firing rates and neurons are sparsely interconnected through excitatory and inhibitory synapses of various strengths (Methods). Following ideas developed in the context of liquid-state (Maass et al., 2002) and echo-state (Jaeger, 2003) models, we assume that this basic network is not designed for any specific

task, but is instead a general purpose dynamical system that will be co-opted for particular applications through subsequent synaptic modification. As a result, the connectivity and synaptic strengths of the network are chosen randomly (Methods). For the parameters we use, the initial state of the network is chaotic (figure 2A).

To specify a task for the networks of figure 1, we must define their outputs. In a full model, this would involve simulating activity all the way out to the periphery. In the absence of such a complete model, we need to have a way of describing what the network is "doing", and here we follow another suggestion from the liquid- and echo-state approach (Maass et al., 2002; Jaeger, 2003; see also Buonomano and Merzenich, 1995). We define the network output through a weighted sum of its activities. Denoting the activities of the network neurons at time t by assembling them into a column vector $\mathbf{r}(t)$ and the weights connecting these neurons to the output by another column vector $\mathbf{w}$, we define the network output as

$$z(t) = \mathbf{w}^{\mathrm{T}}\mathbf{r}(t)\,. \tag{1}$$

Multiple readouts can be defined in a similar manner, each with its own set of weights, but we restrict the discussion to one readout at this point. Although a linear readout is a useful way of defining what we mean by the output of a network, it should be kept in mind that it is a computational stand-in for complex transduction circuitry. For this reason, we refer to the output-generating element as a unit rather than a neuron, and we call the components of $\mathbf{w}$ weights rather than synaptic strengths.

Having specified the network output, we can now define the task we want the network to perform, which is to set $z(t) = f(t)$ for a pre-defined target function $f(t)$. In most of the examples we present, the goal is to make a network produce the target function in the absence of any input. Later, we consider the more conventional network task of generating outputs that depend on inputs to the network in a specified way. Due to stability issues, this is an easier task than generating target functions without inputs, so we mainly focus on the no-input case.

In the initial instantiation of our model (figure 1A), we follow Jaeger and Haas (2004) and modify only the output weight vector $\mathbf{w}$. All other network connections are left unchanged from their initial, randomly chosen values. The critical element that makes such a procedure possible is a feedback loop that carries the output z back into the network (figure 1A). Learning cannot be accomplished in a network receiving no external input without including such a loop. The strengths of the synapses from this loop onto the neurons of the network are chosen randomly and left unmodified. The strength of the feedback synapses is of order 1 whereas that of synapses between neurons of the recurrent network is of order 1 over the square root of the number of recurrent synapses per neuron. The feedback synapses are made stronger so that the feedback pathway has an appreciable effect on the activity of the recurrent network. Later, when we consider the architectures of figures 1B & 1C, we will no longer need such strong synapses.

FORCE Learning

Training in the presence of the feedback loop connecting the output in figure 1A back to the network is challenging because modifying the readout weights produces delayed effects that can be difficult to calculate. Modifying $\mathbf{w}$ has a direct effect on the output z given by equation 1, and it is easy to determine how to change $\mathbf{w}$ to make z closer to f through this direct effect. However, the feedback loop in figure 1A gives rise to a delayed effect when the resulting change in the output caused by modifying $\mathbf{w}$ propagates repeatedly along the feedback pathway and through the network, changing network activities. Because of this delayed effect, a weight modification that at first appears to bring z closer to f may later cause it to deviate away. This problem of delayed effects arises when attempting to modify synapses in any recurrent architecture, including those of figure 1B and 1C.

As stated in the Introduction, Jaeger and Haas (2004) eliminated the problem of delayed effects by clamping feedback during learning. In other words, the output of the network, given by equation 1 was compared with f to determine an error that controlled modification of the readout weights, but this output was not fed back to the network during training. Instead the feedback pathway was clamped to the target function f. The true output was only fed back to the network after training was completed.

We take another approach, which does not require any clamping or manipulation of the feedback pathway, it relies solely on error-based modification of the readout weights. In this scheme, we allow output errors to be fed back into the network, but we keep them small by making rapid and effective weight modifications. As long as output errors are small enough, they can be fed back without disrupting learning, i.e. without introducing significant delayed, reverberating effects. Because the method requires tight control of a small (first-order) error, we call it First-Order, Reduced and Controlled Error or FORCE learning. Although the FORCE procedure holds the feedback signal close to its desired value, it does not completely clamp it. This difference, although numerically small, has extremely significant implications for network stability. Small differences between the actual and desired output of the network during training allow the learning procedure to sample instabilities in the recurrent network and stabilize them.

A learning algorithm suitable for FORCE learning must rapidly reduce the magnitude of the difference between the actual and desired output to a small value, and then keep it small while searching for and eventually finding a set of fixed readout weights that can maintain a small error without further modification. A number of algorithms are capable of doing this (Discussion). All of them involve updates to the values of the weights at times separated by an interval Δt. Each update consists of evaluating the output of the network, determining how far this output deviates from the target function, and modifying the readout weights accordingly. Note that Δt is the interval of time between modifications of the readout weights, not the basic integration time step for the network simulation, which can be smaller than Δt.

At time t, the training procedure starts by sampling the network output, which is given at

this point by $\mathbf{w}^T(t-\Delta t)\mathbf{r}(t)$. The reason that the weights appear here evaluated at time $t-\Delta t$ is that they have not yet been updated by the modification procedure, so they take the same values that they had at the end of the previous update. Comparing this output with the desired target output $f(t)$, we define the error

$$e_-(t) = \mathbf{w}^T(t-\Delta t)\mathbf{r}(t) - f(t)\,. \tag{2}$$

The minus subscript signifies that this is the error prior to the weight update at time t. The next step in the training process is to update the weights from $\mathbf{w}(t-\Delta t)$ to $\mathbf{w}(t)$ in a way that reduces the magnitude of $e_-(t)$. Immediately after the weight update, the output of the network is $\mathbf{w}^T(t)\mathbf{r}(t)$, assuming that the weights are modified rapidly on the scale of network evolution (Discussion). Thus, the error after the weight update is

$$e_+(t) = \mathbf{w}^T(t)\mathbf{r}(t) - f(t)\,, \tag{3}$$

with the plus subscript signifying the error after the weights have been updated.

The goal of any weight modification scheme is to reduce errors by making $|e_+(t)| < |e_-(t)|$ and also to converge to a solution in which the weight vector is no longer changing so that training can be terminated. This latter condition corresponding to making $e_+(t)/e_-(t) \to 1$ by the end of training. In most training procedures, these two conditions are accompanied by a steady reduction in the magnitude of both errors (e_+ and e_-) over time, which are both quite large during the early stages of training. FORCE learning is unusual in that the magnitudes of these errors are small throughout the learning process, although they are similarly reduced over time. This is done by making a large reduction in their size at the time of the first weight update, and then maintaining small errors throughout the training process that decrease with time.

If the training process is initialized at time $t=0$, the first weight update will occur at time Δt. A weight modification rule useful for FORCE learning should make $|e_+(\Delta t)|$, the error after the first weight update has been performed, small, and then keep $|e_-(t)|$ small while slowly increasing $e_+(t)/e_-(t) \to 1$. Given a small magnitude of $e_+(t-\Delta t))$, $e_-(t)$, which is equal to $e_+(t-\Delta t)$ plus a term of order Δt, is kept small by keeping the updating interval Δt sufficiently short. This means that learning can be performed with an error that starts and stays small.

As stated above, several modification rules meet the requirements of FORCE learning, but the recursive least squares (RLS) algorithm is particularly powerful (Haykin, 2002), and we use it here (see Discussion and Supplementary Materials for another, simpler algorithm). In RLS modification,

$$\mathbf{w}(t) = \mathbf{w}(t-\Delta t) - e_-(t)\mathbf{P}(t)\mathbf{r}(t)\,, \tag{4}$$

where $\mathbf{P}(t)$ is an $N \times N$ matrix that is updated at the same time as the weights according to the rule

$$\mathbf{P}(t) = \mathbf{P}(t-\Delta t) - \frac{\mathbf{P}(t-\Delta t)\mathbf{r}(t)\mathbf{r}^T(t)\mathbf{P}(t-\Delta t)}{1+\mathbf{r}^T(t)\mathbf{P}(t-\Delta t)\mathbf{r}(t)}\,. \tag{5}$$

The algorithm also requires an initial value for $\mathbf{P}$, which is taken to be

$$\mathbf{P}(0) = \frac{\mathbf{I}}{\alpha}, \tag{6}$$

where I is the identity matrix and α is a constant parameter. Equation 4 can be viewed as a standard delta-type rule (that is, a rule involving the product of the error and the presynaptic firing rate), but with multiple learning rates given by the matrix $\mathbf{P}$, rather than by a scalar quantity. In this algorithm, $\mathbf{P}$ is a running estimate of the inverse of the correlation matrix of the network rates $\mathbf{r}$ plus a regularization term (Haykin, 2002), i.e. $\mathbf{P} = \left(\sum_t \mathbf{r}(t)\mathbf{r}^{\mathrm{T}}(t) + \alpha\mathbf{I}\right)^{-1}$.

It is straightforward to show that the RLS rule satisfies the conditions necessary for FORCE learning. First, if we assume that the initial readout weights are zero for simplicity (this is not essential), the above equations imply that the error after the first weight update is

$$e_+(\Delta t) = -\frac{\alpha f(\Delta t)}{\alpha + \mathbf{r}^{\mathrm{T}}(\Delta t)\mathbf{r}(\Delta t)}. \tag{7}$$

The quantity $\mathbf{r}^{\mathrm{T}}\mathbf{r}$ is of order N, the number of neurons in the network, so as long as $\alpha \ll N$, this error is small, and its size can be controlled by adjusting α (see below). Furthermore, at subsequent times, the above equations imply that

$$e_+(t) = e_-(t)\left(1 - \mathbf{r}^{\mathrm{T}}(t)\mathbf{P}(t)\mathbf{r}(t)\right), \tag{8}$$

The quantity $\mathbf{r}^{\mathrm{T}}\mathbf{P}\mathbf{r}$ varies over the course of learning from something close to 1 to a value that asymptotically approaches 0, and it is always positive. This means that the size of the error is reduced by the weight update, as required, and ultimately $e_+(t)/e_-(t) \to 1$.

The parameter α, which acts as a learning rate, should be adjusted depending on the particular target function being learned. Small α values result in fast learning, but sometimes make weight changes so rapid that the algorithm becomes unstable. In those cases, larger α should be used (subject to the constraint $\alpha \ll N$), but if α is too large, the FORCE algorithm may not keep the output close to the target function for a long enough time, causing learning to fail. In practice, values from 1 to 100 are effective, depending on the task.

In addition to dealing with feedback, FORCE learning must control the chaotic activity of the network during the training process. In this regard, it is important to note that the network we are considering is being driven through the feedback pathway by a signal approximately equal to the target function. Such an input can induce a transition between chaotic and non-chaotic states (Molgedey et al., 1992; Bertchinger and Natschläger, 2004; Rajan et al., 2008). This is how the problem of chaotic activity can be avoided. Provided that the feedback signal is of sufficient amplitude and frequency to induce a transition to a non-chaotic state (the required properties are discussed in Rajan et al., 2008), learning can take place in the absence of chaotic activity, even though the network is chaotic prior to learning.

Examples of FORCE Learning

Figure 2A-C illustrates how the activity of an initially chaotic network can be modified so that it ends up producing a periodic, triangle-wave output autonomously. Initially, with the output weight vector **w** chosen randomly, the neurons in the network exhibit chaotic spontaneous activity, as does the network output (figure 2A). When we start FORCE learning, the weights of the readout connections begin to fluctuate rapidly, which immediately changes the activity of the network so that it is periodic rather than chaotic and forces the output to match the target triangle wave (figure 2B). The progression of learning can be tracked by monitoring the size of the fluctuations in the readout weights (orange trace in figure 2B), which diminish over time as the learning procedure establishes a set of static weights that generate the target function without requiring modification. At this point, learning can be turned off, and the network continues to generate the triangle wave output on its own indefinitely (figure 2C). The learning process is rapid, converging in only four cycles of the triangle wave in this example.

FORCE learning can be used to modify networks that are initially in a chaotic state so that they autonomously produce a wide variety of outputs (Figure 2D-K). In these examples, training typically converges in about 1000τ, where τ is the basic time constant of the network, which we set to 10 ms. This means learning takes about 10 s of simulated time. Networks can be trained to produce periodic functions of different complexity and form (figures 2D-G and I), even when the target function is corrupted by noise (figure 2F). The dynamic range of the outputs that chaotic networks can be trained to generate by FORCE learning is impressive. For example, a 1000 neuron network with a time constant of 10 ms can produce sine wave outputs with periods ranging from 60 ms to 8 s (figure 2I).

FORCE learning is not restricted to periodic functions. For example, a network can be trained to produce an output matching one of the dynamic variables of the three-dimensional chaotic Lorenz attractor (Methods, see also Jaeger and Haas, 2004), although in this case, because the target is itself a chaotic process, a precise match between output and target can only last for a finite amount of time (figure 2H). After the two traces diverge, the network still produces a trace that looks like it comes from the Lorenz model.

FORCE learning can also produce a segment matching a one-shot, non-repeating target function (figure 2J). To produce such a one-shot sequence, the network must be initialized properly, and we do this by introducing a fixed-point attractor as well as the network configuration that produces the one-shot sequence. This is done by adding a second readout unit to the network that also provides feedback (Methods, network diagram in figure 2J). The first feedback unit induces the fixed point corresponding to a constant z output (horizontal red line in figure 2J), and then the second unit induces the target pattern (red trace between the arrows in figure 2J). As shown below, initialization can also be achieved through appropriate input.

As discussed above, FORCE learning must induce a transition in the network from chaotic to non-chaotic activity during training. This requires an input to the network, through the

feedback loop in our case, of sufficient amplitude. If we try to train a network to generate a target function with too small an amplitude, the activity of the network neurons remains chaotic even after FORCE learning is activated (figure 2K). In this case, learning does not converge to a successful solution. There are a number of solutions to this problem. It is possible for the network to generate low amplitude oscillatory and non-oscillatory functions if these are displaced from zero by a constant shift. Alternatively, the networks shown in figure 1B and C can be trained to generate low amplitude signals centered near zero.

PCA Analysis of FORCE Learning

The activity of a network that has been modified by the FORCE procedure to produce a particular output can be analyzed by principal component analysis (PCA). For a network producing the periodic pattern shown in figure 3A, the distribution of PCA eigenvalues (figure 3C) indicates that the trajectory of network activity lies primarily in a subspace that is of considerably lower dimension than the number of network neurons. The projections of the network activity vector $\mathbf{r}(t)$ onto the PC vectors form a set of orthogonal basis functions (figure 3B) from with the target function is generated. An accurate approximation of the network output (brown trace in Figure 3A) can be generated using the basis functions derived from only the first 8 principal components (with components labeled in decreasing order of the size of their eigenvalues). These 8 components are not the whole story, however, because, along with generating the target function, the network must be stable. If we express the readout weight vector in terms of its projections onto the PC vectors of the network activity, we find that learning sets about the top 50 of these projections to uniquely specified values (figure 3E). The remaining projections take different values from one learning trial to the next, depending on initial conditions (figure 3F). This multiplicity of solutions greatly simplifies the task of finding successful readout weights.

The uneven distribution of eigenvalues shown in figure 3C illustrates why the RLS algorithm works so well for FORCE learning. As mentioned previously, the matrix $\mathbf{P}$ acts as a set of learning rates for the RLS algorithm. This is seen most clearly by shifting to a basis in which $\mathbf{P}$ is diagonal. Assuming learning has progressed long enough for $\mathbf{P}$ to have converged to the inverse correlation matrix of $\mathbf{r}$, the diagonal basis is achieved by projecting $\mathbf{w}$ and $\mathbf{r}$ onto the PC vectors. Doing this, it is straightforward to show that the learning rate for the component of $\mathbf{w}$ aligned with PC vector a after M weight updates is $1/(M\lambda_a + \alpha)$, where λ_a is the corresponding PC eigenvalue. This rate divides the RLS process into two stages, one when $M < \alpha/\lambda_a$ in which the major role of weight modification is to control the output (set it close to f), and another when $M > \alpha/\lambda_a$ in which the goal is learning, that is, finding a static weight that accomplishes the task. Components of $\mathbf{w}$ with large eigenvalues quickly enter the learning phase, whereas those with small eigenvalues spend more time in the control phase (figure 3D). Controlling components with small eigenvalues allows weight projections in dimensions with large eigenvalues to be learned.

The learning rate for all components during the control phase is $1/\alpha$. During the learning phase, the rate for PC component a is proportional to $1/\lambda_a$. The average rate of change (as

opposed to just the learning rate) of the projection of the output weight vector onto principal component a is proportional to $\sqrt{\lambda_a}/(M\lambda_a + \alpha)$ because the factor of $\mathbf{r}$ in equation 4 introduces a term proportional to $\sqrt{\lambda_a}$, so the full rate of change for large M goes as $1/\sqrt{\lambda_a}$. This is exactly what it should be, because in the expression for z, this change is multiplied by the projection of $\mathbf{r}$ onto PC vector a, which again has an amplitude proportional to $\sqrt{\lambda_a}$. Thus, RLS, by having rates of change of $\mathbf{w}$ proportional to $1/\sqrt{\lambda_a}$ in the PC basis, allows all the projections to, potentially, contribute equally to the output of the network.

Comparison of Echo-State and FORCE Feedback

In echo-state learning (Jaeger and Haas, 2004), the feedback signal during training was set equal to the target function $f(t)$. In FORCE learning, the feedback signal is $z(t)$ during training. To compare these two methods, we introduce a mixed feedback signal, setting the feedback equal to $\gamma f(t) + (1-\gamma)z(t)$ during training. Thus, $\gamma = 0$ corresponds to FORCE learning and $\gamma = 1$ to echo-state learning, with intermediate values interpolating between these two approaches.

Training to produce the output of figure 3A, we find the network is only stable on the majority of trials when $\gamma < 0.15$, in other words close to the FORCE limit (figure 4A). Furthermore, in this γ range, the error in the output after training increases as a function of γ, meaning $\gamma = 0$ performs best (figure 4B). For a typical instability of pure echo-state learning, the output matches the target briefly after learning is terminated, but then it deviates away (figure 4C). Because this stability problem arises from the failure of the network to sample feedback fluctuations, it can be alleviated somewhat by introducing noise into the feedback loop during training (Jaeger and Haas, 2004, introduced noise into the network, which is less effective). Doing this, we find that pure echo-state learning converges on about 50% of the trials, but the error on these is significantly larger than for pure FORCE learning.

Advantages of Chaotic Spontaneous Activity

To study the effect of spontaneous chaotic activity on network performance, we introduce a factor g that scales the strengths of the recurrent connections within the network. Networks with $g < 1$ are inactive prior to training, whereas networks with $g > 1$, exhibit chaotic spontaneous activity (Sompolinsky et al., 1988) that gets more irregular and fluctuates more rapidly as g is increased beyond 1 (we typically use $g = 1.5$).

The number of cycles required to train a network to generate the periodic target function shown in figure 3A drops dramatically as a function of g, continuing to fall as g gets larger than 1 (figure 5A). The average root-mean-square (RMS) error, indicating the difference between the target function and the output of the network after FORCE learning, also decreases with g (figure 5B). Another measure of training success is the magnitude of the readout weight vector $|\mathbf{w}|$ (figure 5C). Large values of $|\mathbf{w}|$ indicate that the solution found by a learning process involves cancellations between large positive and negative contribu-

tions. Such solutions tend to be unstable and sensitive to noise. The magnitude of the weight vector falls as a function of g and takes its smallest values in the region $g > 1$ characterized by chaotic spontaneous activity.

These results indicate that networks that are initially in a chaotic state are quicker to train and produce more accurate and robust outputs than non-chaotic networks. Learning works best when $g > 1$ and, in fact, fails in this example for networks with $g < 0.75$. This might suggest that the larger the g value the better, but there is an upper limit. Recall that FORCE learning does not work if the feedback from the readout unit to the network fails to suppress the chaos in the network. For any given target function and set of feedback synaptic strengths, there is an upper limit for g beyond which chaos cannot be suppressed by FORCE learning. Indeed, the range of g values in figure 5 terminates at $g = 1.56$ because learning did not converge for higher g values due to this problem. Thus, the best value of g for a particular target function is at the "edge of chaos" (Bertchinger and Natschläger, 2004), that is the g value just below the point where FORCE learning fails to suppress chaotic activity during training.

Distorted and Delayed Feedback

The linear readout unit was introduced into the network model as a stand-in for a more complex, un-modeled peripheral system, in order to define the output of the network. The critical information provided by the readout unit is the error signal needed to guide weight modification, so its biological interpretation should be as a system that computes or estimates the deviation between an action generated by a network and the desired action. However, in the network configuration presented to this point (figure 1A), the readout unit, in addition to generating the error signal that guides learning, is also the source of feedback. Given that the output in a biological system is actually the result of a large amount of nonlinear processing and that feedback, whether proprioceptive or a motor efference copy, may have to travel a significant distance before returning to the network, we begin this section by examining the effect of introducing delays and nonlinear distortions along the feedback pathway from the readout unit to the network neurons.

The FORCE learning scheme is remarkably robust to distortions introduced along the feedback pathway (figure 6A). Nonlinear distortions of the feedback signal can be introduced as long as they do not diminish the temporal fluctuations of the output to the point where chaos cannot be suppressed. We have also introduced low-pass filtering into the feedback pathway, which can be quite extreme before the network fails to learn. Delays can be more problematic if they are too long. The critical point is that FORCE learning works as long as the feedback is of an appropriate form to suppress the initial chaos in the network. This means that the feedback really only has to match the period or the duration of the target function and roughly have the same frequency content.

FORCE Learning with Other Network Architectures

Even allowing for distortion and delay, the feedback pathway, originating as it does from

the linear readout unit, is a non-biological element of the network architecture of figure 1A. To address this problem, we consider two ways of separating the feedback pathway from the linear readout of the network and modeling it more realistically. The first is to provide feedback to the network through a second neural network (figure 1B) rather than via the readout unit. To distinguish the two networks, we call the original network, present in figure 1A, the generator network and this new network the feedback network. The feedback network has nonlinear, dynamic neurons identical to those of the generator network, and is recurrently connected. Each unit of the feedback network produces a distinct output that is fed back to a subset of neurons in the generator network, so the task of carrying feedback is shared across multiple neurons. This repairs two biologically implausible aspects of the architecture of figure 1A: the strong feedback synapses mentioned above and the fact that every neuron in the network receives the same feedback signal.

When we include a feedback network (figure 1B), FORCE learning takes place both on the weights connecting the generator network to the readout unit (as in the architecture of figure 1A) and on the synapses connecting the generator network to the feedback network (red connections in figure 1B). Separating feedback from output introduces a credit-assignment problem because changes to the synapses connecting the generator network to the feedback network do not have a direct effect on the output. To solve this problem within the FORCE learning scheme, we treat every neuron subject to synaptic modification as if it were the readout unit, even when it is not. In other words, we apply equations 4 & 5 to every synapse connecting the generator network to the feedback network (Supplementary Materials), and we also apply them to the weights driving the readout unit. When we modify synapses onto a particular neuron of the feedback network, the vector $\mathbf{r}$ in these equations is composed of the firing rates of generator network neurons presynaptic to that feedback neuron, and the weight vector $\mathbf{w}$ is replaced by the strengths of the synapses it receives from these presynaptic neurons. However, the same error term that originates from the readout (equation 2) is used in these equations whether they are applied to the weights of the readout unit or synapses onto neurons of the feedback network (Methods). The form of FORCE learning we are using is cell autonomous, so no communication of learning-related information between neurons is required to implement these modifications, except that they all use a global error signal.

FORCE learning with a feedback network and independent readout unit can generate complex outputs similar to those in figure 4, although parameters such as α (equation 6) may require more careful adjustment. After training, when the output of these networks matches the target function, the activities of neurons in the feedback network do not, despite the fact that their synapses are modified by the same algorithm as the readout weights (figure 6B). This difference is due to the fact that the feedback network neurons receive input from each other as well as from the generator network, and these other inputs are not modified by the FORCE procedure. Differences between the activity of feedback network neurons and the output of the readout unit can also arise from different values of the synapses and the readout weights prior to learning.

With a separate feedback network, the feedback to an individual neuron of the generator network is a random combination of the activities of a subset of feedback neurons, summed through random synaptic weights. While these sums bear a certain resemblance to the target function, they are not identical to it nor are they identical for different neurons of the generator network. Nevertheless, FORCE learning works. This extends the result of figure 6A, showing not only that the feedback does not have to be identical to the network output, but that it does not even have to be identical for each neuron of the generator network.

Why does this form of learning, in which every neuron with synapses being modified is treated as if it were producing the output, work? In the example of figure 6B, the connections from the generator network to the readout unit and to the feedback network are sparse and random (Methods), so that neurons in the feedback network do not receive the same inputs from the generator network as the readout unit. However, suppose for a moment that each neuron of the feedback network, as well as the readout unit, received synapses from all of the neurons of the generator network. In this case, the changes to the synapses onto the feedback neurons would be identical to the changes of the weights onto the readout unit and therefore would induce a signal identical to the output into each neuron of the feedback network. This occurs, even though there is no direct connection between these two circuit elements, because the same learning rule with the same global error is being applied in both cases.

The explanation of why FORCE learning works in the feedback network when the connections from the generator network are sparse rather than all-to-all (as in figure 6B) relies on the accuracy of randomly sampling a large system (Sussillo, 2009). With sparse connectivity, each neuron samples a subset of the activities within the full generator network, but if this sample is large enough, it can provide an accurate representation of the leading principal components of the activity of the generator network that drive learning. This is enough information to allow learning to proceed. For figure 6B, we used an extremely sparse connectivity (Methods) to illustrate that FORCE learning can work even when the connections of the units being modified are highly non-overlapping.

The original generator network we used (figure 1A) is recurrent and can produce its own feedback. This means that we should be able to apply FORCE learning to the synapses of the generator network itself, in the arrangement shown in figure 1C. To implement FORCE learning within the generator network (Supplementary Materials), we modify every synapse in that network using equations 4-5. To apply these equations, the vector $\mathbf{w}$ is replaced by the set of synapses onto a particular neuron being modified, and $\mathbf{r}$ is replaced by the vector formed from the firing rates of all the neurons presynaptic to that network neuron. As in the example of learning in the feedback network, FORCE learning is also applied to the readout weights, and the same error, given by equation 2, is used for every synapse or weight being modified.

FORCE learning within the network can produce a complex target output (figure 6C). An argument similar to that given for learning within the feedback network can be applied

to FORCE learning for synapses within the generator network. To illustrate how FORCE learning works, we express the total current into each neuron of the generator network as the sum of two terms. One is the current produced by the original synaptic strengths prior to learning, $\sum_j J_{ij}(0)r_j(t)$ for neuron i. The other is the extra current generated by the learning-induced changes in these synapses, $\sum_j \left(J_{ij}(t) - J_{ij}(0)\right) r_j(t)$. The first term, as well as the total current, is different for each neuron of the generator network because of the random initial values of the synaptic strengths. The second, learning-induced current, however, is virtually identical to the target function for each neuron of the network (figure 6C, cyan). Thus, FORCE learning induces a signal representing the target function into the network, just as it does for the architecture of figure 1A, but in a subtler and more biologically realistic manner.

Output patterns like those in figure 2 can be reproduced by FORCE learning applied within the generator or feedback networks. In the following sections, we illustrate the capacity of these forms of FORCE learning while, at the same time, introducing new tasks. All of the examples shown can be reproduced using all three of the architectures in figure 1, but for compactness we show results from learning in the generator network in figure 7 and learning in the feedback network in figure 8. For the interested reader, Matlab files that implement FORCE learning in the different architectures are included with the Supplementary Materials.

Switching Between Multiple Outputs and Input-Output Mapping with Memory

The examples to this point have involved a single target function. We can train networks with the architecture of figure 1C in both sparse and fully connected configurations (we illustrate the sparse case) to produce multiple functions, with a set of inputs controlling which is generated at any particular time. We do this by introducing static control inputs to the network neurons (figure 7A) and pairing each desired output function with a particular input pattern (Methods). The constant values of the control inputs are chosen randomly. When a particular target function is being either trained or generated, the control inputs to the network are set to the corresponding static pattern and held constant until a different output is desired. The control inputs do not supply any temporal information to the network, they act solely as a switching signal to select a particular output function. The result is a single network that can produce a number of different outputs depending on the values of the control inputs (figure 7B).

Up to now, we have treated the network we are studying as a source of what are analogous to motor output patterns. Networks can also generate complex input/output maps when inputs are present. Figure 7C shows a particularly complex example of a network that functions as a 4-bit memory that is robust to input noise. This network has 8 inputs that randomly connect to neurons in the network and are functionally divided into pairs (Methods). The input values are held at zero except for short pulses to positive values that act as ON and OFF commands for the 4 readout units. Input 1 is the ON command for output 1 and input 2 is its OFF command. Similarly, inputs 3 and 4 are the ON and OFF commands for output

2, and so on. Turning on an output means inducing a transition to a state with a fixed positive value of 1, and turning it off means switching it to -1. After FORCE learning, the inputs correctly turn the appropriate outputs on and off with little crosstalk between inputs and inappropriate outputs (figure 7C). This occurs despite the random connectivity of the network, which means that the inputs do not segregate into different channels. This example requires the network to have, after learning, 16 different fixed point attractors, one for each of the 4^2 possible combinations of the 4 outputs, and the correct transitions between these attractors induced by pulsing the 8 inputs.

A Motion Capture Example

Finally, we consider an example of running and walking based on data obtained from human subjects performing these actions while wearing a suit that allows variables such as joint angles to be measured (also studied by Taylor et al., 2008 using a different type of network and learning procedure). These data, from the CMU Motion Capture Library, consist of 95 joint angles measured over hundreds of time steps.

We implemented this example using all the architectures in figure 1 in both sparse and fully connected configurations with similar results (we show a sparse example using the architecture of figure 1B). Producing all 95 joint angle sequences in the data sets requires that these networks have 95 readout units. For internal learning, subsets of neurons subjected to learning were assigned to each readout unit and trained using the error generated by that unit (Methods). Although running and walking might appear to be periodic motions, in fact the joint angles in the real data are non-periodic. For this reason, we introduced static control inputs to initialize the network prior to starting the running or walking motion. Because we wanted a single network to generate both motions, we also used the control inputs to switch between them, as in figure 7A. The successfully trained networks produced both motions (figure 8; for an animated demo showing all the architectures of figure 1 see the avi files included with the Supplementary Materials) demonstrating that a single chaotic recurrent network can generate multiple, high-dimensional, non-periodic patterns that resemble complex human motions.

Discussion

In the Introduction, we mentioned that FORCE learning could be viewed either as a model for learning-induced modification of biological networks or, more simply, as a method for constructing models of these networks. Our results should be evaluated in light of both of these interpretations.

FORCE learning solves some, but certainly not all, of the problems associated with applying ideas about learning from mathematical neural networks to biological systems. Biological networks exhibit complex and irregular spontaneous activity that probably has both chaotic and stochastic sources. FORCE learning provides an approach to network training that can be applied under these conditions (see, in particular, figure 2F). Furthermore, training does not require any reconfiguration of the network or changes in its dynamics

other than the introduction of synaptic modification. Finally, the networks constructed by FORCE learning are more stable than in previous approaches.

FORCE learning relies on an error signal that, in our examples, is based on a readout unit that is not intended to be a realistic circuit element. It is not clear how the error is computed in biological systems. This is a problem for all models of supervised learning. In motor learning, we imagine that the target function is generated by an internal model of a desired movement and that circuitry exists for comparing this internal model with the motor signal generated by the network and for producing a modulatory signal that guides synaptic plasticity. The cerebellum has been proposed as a possible locus for such internal modeling (Miall et al., 1993). Examples like that of figure 8, which involve multiple outputs, require multiple error signals. For figure 8, we subdivided the network being trained into different regions in which plasticity was controlled by a different error signal. If the error is carried by a neuromodulator, this would require multiple pathways (though not necessarily multiple modulators) with at least some spatial targeting. If the error signal is transmitted as in the case of the climbing fibers of the cerebellum, multiple error signals are more straightforward to handle. Examples with a single output only require a single global error signal.

It is also not known how the error signal, once generated, controls synaptic plasticity. Again, this is a problem associated with all models of error- or reward-based learning. FORCE learning adds the condition that this modification act rather quickly compared to the timescale of the action being learned, at least during the initial phases of learning. Both because it is under the control of an error signal and because it acts rapidly, the plasticity required does not match that of typical long-term potentiation experiments, and it is a challenge raised by this work to uncover how such rapid plasticity can be realized biologically, or if it is realized at all. Whatever the plasticity mechanism, a key component of FORCE learning is producing the roughly correct output even during the initial stages of training. Analogously, people cannot learn fine manual skills by randomly flailing their arms about and having their movement errors slowly diminish over time, which would be analogous to more conventional network learning schemes. FORCE learning reminds us that motor learning works best when the desired motor action is duplicated as accurately as possible during training.

The RLS algorithm we have used is neuron-specific but not synapse-specific. By this we mean that the algorithm uses information about all the inputs to a given neuron to guide modification of its individual synapses. The algorithm requires some fairly involved calculations, although not matrix inversion. It is possible to use a simpler, synapse-specific weight modification procedure in which the matrix $\mathbf{P}$ is replaced by a single learning rate (Supplementary Material). Provided that this scalar rate is adapted over time, FORCE learning can work with such a simpler plasticity mechanism. Nevertheless, RLS is clearly a more powerful algorithm because it adapts the learning rate to the magnitude of different principal components of the network activity. It is possible that a scheme that is simpler and more biologically plausible than RLS can be devised that retains this desirable feature.

The architectures of figures 1B and 1C, where learning occurs within feedback or generator networks, match biological circuits better than that of figure 1A, where feedback comes directly from the readout unit. A key feature of learning in these cases is that network plasticity is accompanied by plasticity along the output or error-computing pathway. Plasticity in multiple areas (at least two, in these examples) coupled by a common error signal is a basic prediction of the model. It is a curious feature that performance is comparable for all three architectures in figure 1, despite that fact that the case of figure 1C involves changing many more synaptic strengths. We do not currently know whether changing synapses within a network offers advantages for the function-generation task. It may, but the modification algorithms developed thus far are not powerful enough exploit these advantages.

We now come to an analysis of FORCE learning as a model-building scheme. We have studied how spontaneously active neural networks can be modified to generate desired outputs, and how control inputs can be used to initiate and select among those outputs. Although this has most direct application to motor systems, it can be generalized to a broader picture of cognitive processing (Yuste et al., 2005; Buonomano and Maass, 2009), as our example of a 4-bit, input-controlled memory suggests.

Ganguli et al. (2008) have discussed the advantages of using an effective delay-line architecture in applications of networks to memory retention. Provided that a feedback loop from the output, as in figure 1A, is in place, a delay line structure within the generator network should be quite effective for function generation as well. However, because we were interested in networks that generate spontaneous activity even in the absence of the output feedback loop, we did not consider such an arrangement in any detail.

The two-step process by which we induced a chaotic network to produce a non-periodic sequence (figures 2J and 8) may have an analog in motor and premotor cortex. The brief fixed point that we introduced to terminate chaotic activity results in a sharp drop in the fluctuations of network neurons just before the learned sequence is generated. Churchland et al. (2006) and Churchland and Shenoy (2007a) have reported just such a drop in variability in their recordings from motor and premotor areas in monkeys immediately before they performed a reaching movement. Except in the simplest of examples, the activity of the generator neurons bears little relationship to the output of the network. Trying to link single-neuron responses to motor actions may thus be misguided. Instead, results from our network models suggest that it may be more instructive to study network-wide modes or patterns of activity extracted by principal components analysis of multi-unit recordings (Fetz, 1992; Robinson, 1992; Churchland and Shenoy, 2007b).

There are some interesting and perhaps unsettling aspects of the networks we have studied. First, the connectivity of the generator network in the architectures of figure 1A and 1B is completely random, even after the network has been trained to perform a specific task. If would be impossible to understand how the generator network "works" by analyzing its synaptic connectivity. Even when the synapses of the generator network are modified (as in figure 1C), there is no obvious relationship between the task being performed and

the connectivity, which is in any case not unique. The lesson here is that the activity, response properties and function of locally connected neurons can be drastically modified by feedback loops passing through distal networks. Circuits may need to be studied with an eye toward how they modulate each other, rather than how they function in isolation.

The architecture of figure 1B, involving a separate feedback network (basal ganglia or cerebellum), may be a way to keep plasticity from disrupting the generator network (motor cortex), a disruption that would be disastrous for all motor output, not merely the current task being learned. Modification of synapses in a second network (as in figure 1B) may dominate when a motor task is first learned, whereas changes within motor cortex (analogous to learning within the network of figure 1C) may be reserved for "virtuoso" highly trained motor actions. Our examples show the power of adding feedback loops as a way of modifying network activity. Nervous systems often seem to be composed of loops within loops within loops. Because adding a feedback loop leaves the original circuit unchanged, this is a non-destructive yet highly flexible way of increasing a behavioral repertoire through learning, as well as during development and evolution.

Methods

All the networks we use are based on firing-rate descriptions of neural activity. To encompass all the models, we write the network equations for the generator network as (note that, in the Results, we called the parameter labeled here as g_{GG} simply g)

$$\tau \frac{dx_i}{dt} = -x_i + g_{GG} \sum_{j=1}^{N_G} J_{ij}^{GG} r_j + g_{Gz} J_i^{Gz} z + g_{GF} \sum_{a=1}^{N_F} J_{ia}^{GF} s_a + \sum_{\mu=1}^{N_I} J_{i\mu}^{GI} I_\mu$$

for $i = 1, 2, \ldots, N_G$ with firing rates $r_i = \tanh(x_i)$. For the feedback network,

$$\tau \frac{dy_a}{dt} = -y_a + g_{FF} \sum_{b=1}^{N_F} J_{ab}^{FF} s_b + g_{FG} \sum_{i=1}^{N_G} J_{ai}^{FG} r_i + \sum_{\mu=1}^{N_I} J_{a\mu}^{FI} I_\mu$$

for $a = 1, 2, \ldots, N_F$ with firing rates $s_a = \tanh(y_a)$. Equation 1 determines z and $\tau = 10$ ms. Sometimes we assign a sparseness p_z to the readout unit, meaning that a random fraction $1 - p_z$ of the components of $\mathbf{w}$ are set and held to zero. The connection matrices are also assigned sparseness parameters, p, meaning that each element is set and held to 0 with probability $1 - p$. Nonzero elements of J^{GG}, J^{GF}, J^{FG}, J^{FF} are drawn independently from Gaussian distributions with zero means and variances equal to the inverses of $p_{GG}N_G$, $p_{GF}N_F$, $p_{FG}N_G$ and $p_{FF}N_F$, respectively. Rows of J^{GI} and J^{FI} have a single non-zero element drawn from a Gaussian distribution with zero mean and unit variance. Elements of J^{Gz} are drawn from a uniform distribution between -1 and 1. Nonzero elements of $\mathbf{w}$ are set initially either to zero or to values generated by a Gaussian distribution with zero mean and variance $1/(p_z N)$.

For figures 2-5 and 6A: $N_G = 1000$, $p_{GG} = 0.1$, $p_z = 1$, $g_{Gz} = 1$, $g_{GF} = 0$, $\alpha = 1.0$, and $N_I = 0$. For figure 5, g_{GG} was varied, otherwise $g_{GG} = 1.5$. For figure 2H, the standard

Lorenz attractor model (see Strogatz, 1994) was used with $\sigma = 10$, $\beta = 8/3$, and $\rho = 28$. The target function was what is conventionally labeled as x divided by 10, to fit it roughly into the range of -1 to 1. For figure 2J, the two readouts and feedback loops are similar except for different random choices for the strengths of the feedback synapses onto the network neurons. The additional readout unit takes two possible states, one called active in which its output is determined by equation 1, and another called inactive in which its output is 0. For further discussion of training in this case, see the Supplementary Materials.

For figures 6B: $N_G = 20000$, $N_F = 95$, $p_{GG} = 0.1$, $p_{GF} = 0.25$, $p_{FG} = 0.025$, $p_{FF} = 0.25$, $p_z = 0.025$, $g_{GG} = 1.5$, $g_{GF} = 1$, $g_{FG} = 1$, $g_{FF} = 1.2$, $\alpha = 1.0$, and $N_I = 0$. RLS modification was applied to $\mathbf{w}$ and J^{FG}.

For figures 6C: $N_G = 750$, $p_{GG} = 0.5$, $p_z = 0.5$, $g_{GG} = 1.5$, $g_{GF} = 0$, $\alpha = 1.0$, and $N_I = 0$. RLS modification was applied to $\mathbf{w}$ and J^{GG}.

For figure 7B: $N_G = 1200$, $p_{GG} = 0.8$, $p_z = 0.8$, $g_{GG} = 1.5$, $g_{GF} = 0$, $\alpha = 80$, and $N_I = 100$. The inputs I_μ where chosen randomly and uniformly over the range -2 to 2 for inputs generating initialization fixed points, and -0.5 to 0.5 for inputs control the choice of output function. RLS modification was applied to $\mathbf{w}$ and J^{GG}, but of the 1200 network neurons, 800 were subject to synaptic modification of their incoming synapses (due to memory considerations).

For figure 7D: $N_G = 1200$, $p_{GG} = 0.8$, $p_z = 1$, $g_{GG} = 1$, $g_{GF} = 0$, $\alpha = 40$, and $N_I = 8$. The elements of the control input vector I_μ had OFF values of 0.0 and ON values of 0.375. RLS modification was applied to $\mathbf{w}$ and J^{GG}, with 800 of the network neurons subject to synaptic modification of their incoming synapses.

For figure 8: Although all network variants in figure 1 were implemented successfully, the following parameters are for the generator-feedback architecture: $N_G = 5000$, $N_F = 285$, $p_{GG} = 0.05$, $p_{GF} = 0.5$, $p_{FG} = 0.185$, $p_{FF} = 0.5$, $p_z = 0.185$, $g_{GG} = 1.5$, $g_{GF} = 2.0$, $g_{FG} = 1.0$, $g_{FF} = 1.5$, $\alpha = 2.0$, and $N_I = 50$. RLS modification was applied to $\mathbf{w}$ and J^{FG}. Motion capture data were downloaded from the Carnegie Mellon University Motion Capture Library (MOCAP) (http://mocap.cs.cmu.edu/). Data set 09_02.amc was used for the running example and data set 08_01.amc for the walking case. The data were preprocessed by a simple moving average filter to remove discontinuities and then interpolated to fill in to 10 times density, which works better for our continuous time models. The resulting joint angles were transformed into exponential form (see Taylor et al., 2006) and the means were removed. Movement through space was ignored, so we modeled a runner or walker on a treadmill. Four sets of control inputs were used, one each for running and walking and two for initial-value fixed points for these motions. Fixed-point inputs were chosen randomly and uniformly over the range -2 to 2 and control inputs over -0.25 to 0.25.

Acknowledgments

Research supported by an NIH Director's Pioneer Award, part of the NIH Roadmap for

Medical Research, through grant number 5-DP1-OD114-02 and by National Institute of Mental Health grant MH-58754. We thank Taro Toyoizumi, Surya Ganguli, Greg Wayne, and Graham Taylor for helpful comments and suggestions.

References

Abarbanel, H.D., Creveling, D.R., and Jeanne, J.M. (2008). Estimation of parameters in nonlinear systems using balanced synchronization. Phys. Rev. E Stat. Nonlin. Soft. Matter Phys. 77, 016208.

Atiya, A.F. and Parlos, A.G. (2000) New results on recurrent network training: Unifying the algorithms and accelerating convergence. IEEE Transactions on Neural Networks 11, 697-709.

Amit, D.J., Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex. 7, 237-252.

Bertchinger, N., and Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. Neural Comput. 16, 1413-1436.

Brunel, N. (2000). Dynamics of networks of randomly connected excitatory and inhibitory spiking neurons. J. Physiol. Paris 94, 445-463.

Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. Nat. Rev. Neurosci. 10, 113-125.

Buonomano D.V., and Merzenich M.M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. Science 267, 1028-1030.

Churchland, M.M., and Shenoy, K.V. (2007a). Delay of movement caused by disruption of cortical preparatory activity. J. Neurophysiol. 9, 348-359.

Churchland, M.M., and Shenoy, K.V. (2007b). Temporal Complexity and Heterogeneity of Single-Neuron Activity in Premotor and Motor Cortex. J. Neurophysiol. 97, 4235-4257

Churchland, M.M., Yu, B.M., Ryu, S.I., Santhanam, G., and Shenoy, K.V. (2006). Neural variability in premotor cortex provides a signature of motor preparation. J. Neurosci. 26, 3697-3712.

Doya, K. (1992). Bifurcations in the learning of recurrent neural networks. In IEEE International Symposium on Circuits and Systems 6, 2777-2780.

Fetz, E. (1992). Are movement parameters recognizably coded in the activity of single neurons? Behavioral and brain sciences 15, 679-690.

Ganguli, S., Huh, D., and Sompolinsky, H. (2008) Memory traces in dynamical systems. Proc . Natl. Acad. Sci. USA 105, 18970-18975.

Haykin, S. (2002). Adaptive Filter Theory. Prentice Hall, Upper Saddle River NJ.

Hinton, G.E., Osindero, S., and Teh, Y.W. (2006). A fast learning algorithm for deep belief nets. Neural Comput. 18, 1527-1554.

Jaeger, H. (2003). Adaptive nonlinear system identification with echo state networks. In Advances in Neural Information Processing Systems 15, S. Becker, S. Thrun, and K. Obermayer, eds. MIT Press, Cambridge MA. pp. 593-600.

Jaeger H., and Haas H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science 304, 78-80.

Maass, W., Joshi, P., and Sontag, E.D. (2007). Computational aspects of feedback in neural circuits. PLoS Comput, Biol, 3, e165.

Maass, W., Matschlager, T., and Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. Neural Comput. 14, 2531-2560.

Miall, R.C., Weir, D.J., Wolpert, D.M., and Stein, J.F. (1993). Is the cerebellum a smith predictor? J. Motor Behav. 25, 203-216.

Molgedey, L., Schuchhardt, J., and Schuster, H.G. (1992). Suppressing chaos in neural networks by noise. Phys. Rev. Lett. 69, 3717-3719.

Pearlmutter, B. (1989). Learning state space trajectories in recurrent neural networks. Neural Comput. 1, 263-269.

Rajan, K., Abbott, L.F., and Sompolinsky, H. (2008). Stimulus-Dependent Suppression of Intrinsic Variability in Recurrent Neural Networks. (submitted).

Robinson, D. (1992). Implications of neural networks for how we think about brain function. Behavioral and Brain Sciences 15, 644-655.

Rumelhart, D.E., Hinton, G.E., and Williams, R.J. (1986). Learning internal representations by error propagation. In Rumelhart, D.E., and McClelland, J.L., eds. (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1, Foundations. MIT Press, Cambridge MA, chapter 8.

Rumelhart, D.E., and McClelland, J.L., eds. (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Vol. 1, Foundations. MIT Press, Cambridge MA.

Sompolinsky, H., Crisanti, A., and Sommers, H.J. (1988). Chaos in Random Neural Networks. Phys. Rev. Lett. 61, 259-262.

Sussillo, D. (2009) Learning in chaotic recurrent networks. Columbia University Ph.D. Thesis. pp. 93-102.

Taylor, G.W., Hinton, G.E., and Roweis, S. (2006). Modeling human motion using binary latent variables. Advances in Neural Information Processing Systems. MIT Press, Cambridge MA.

van Vreeswijk, C., and Sompolinsky, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274, 1724-1726.

Williams, R.J., and Zipser, D. (1989). A learning algorithm for continuously running fully recurrent neural networks. Neural Comput. 1, 270-280.

Yuste R., MacLean, J.N., Smith, J., and Lansner A. (2005). The cortex as a central pattern generator. Nature Rev. Neurosci. 6, 477-483.

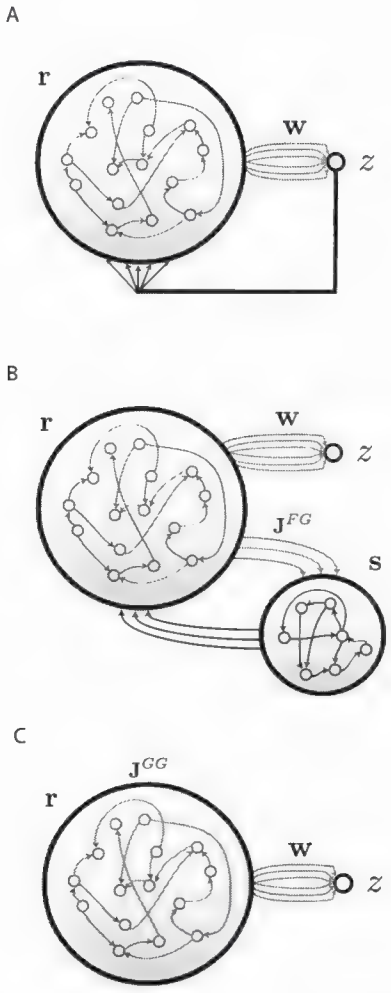

Figure 1: Network architectures. In all three cases, a recurrent generator network with firing rates $\mathbf{r}$ drives a linear readout unit with output z through weights $\mathbf{w}$ (red) that are modified during training. Only connections shown in red are subject to modification. A) Feedback to the generator network (large network circle) is provided by the readout unit. B) Feedback to the generator network is provided by a separate feedback network (smaller network circle). Neurons of the feedback network are recurrently connected and receive input from the generator network through synapses of strength $\mathbf{J}^{\mathrm{FG}}$ (red), which are modified during training. C) A network with no external feedback. Instead, feedback is generated within the network and modified by applying FORCE learning to the synapses with strengths $\mathbf{J}^{\mathrm{GG}}$ internal to the network (red).

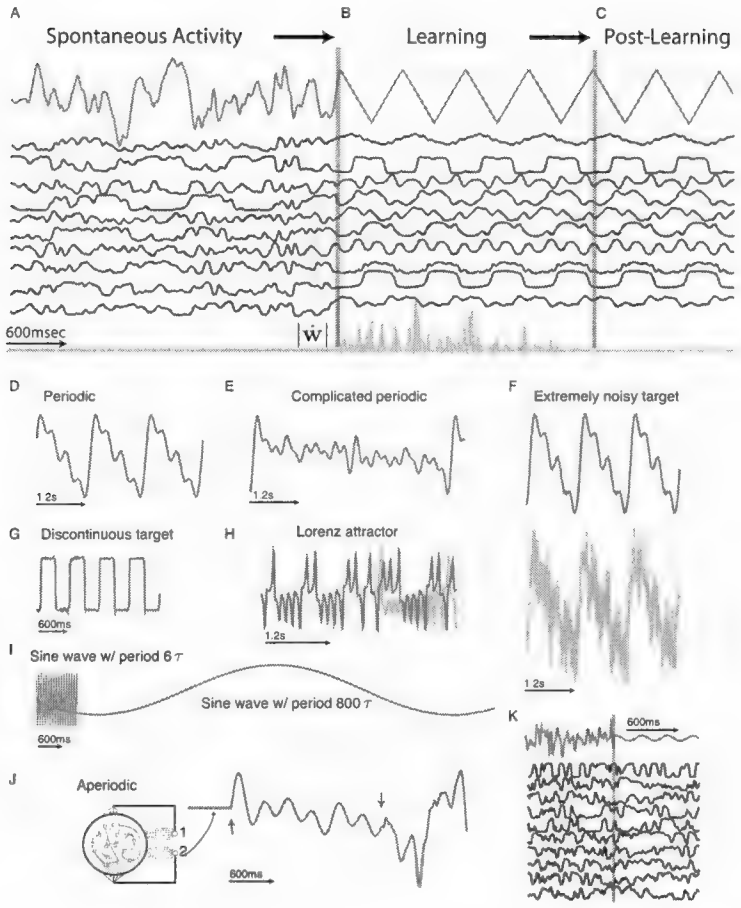

Figure 2: FORCE learning in the network of figure 1A. A-C) The FORCE training sequence. Network output, z, is in red, the firing rates of 10 sample neurons from the network are in blue and the orange trace is the magnitude of the time derivative of the readout weight vector. A) Before learning, network activity and output are chaotic. B) During learning, the output matches the target function, in this case a triangle wave and the network activity is periodic because the readout weights fluctuate rapidly. These fluctuations subside as learning progresses. C) After training, the network activity is periodic and the output matches the target without requiring any weight modification. D-K) Examples of FORCE Learning. Red traces are network outputs after training with the network running autonomously. Green traces, where not covered by the matching red traces, are target functions. D) Periodic function composed of 4 sinusoids. E) Periodic function composed of 16 sinusoids. F) Periodic function of 4 sinusoids learned from a noisy target function. G) Square-wave. H) The Lorenz attractor. Initial conditions of the network and the target were matched at the beginning of the traces. I) Sine waves with periods of 60 ms and 8 s. J) A one-shot example using a network with two readout units (circuit insert). When activated, unit 1 creates the fixed point to the left of the left-most blue arrow, establishing the appropriate initial condition. Unit 2 then produces the sequence between the two blue arrows. When the sequence is concluded, the network output returns to being chaotic. K) A low amplitude sine wave (right of gray line) for which the FORCE procedure does not control network chaos (blue traces) and learning fails.

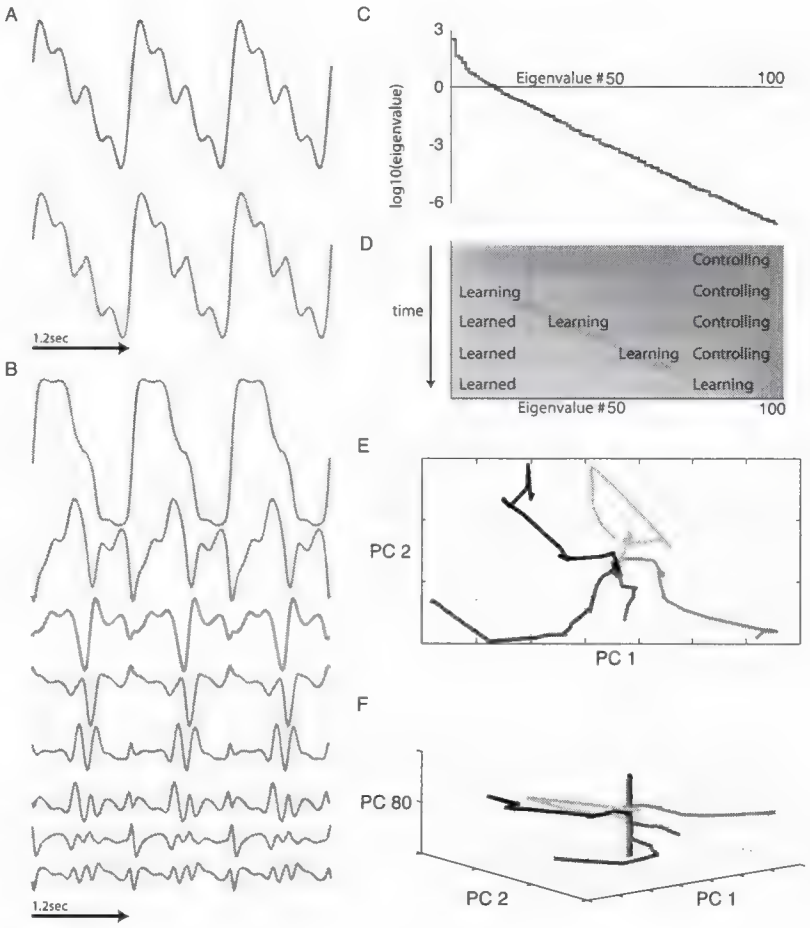

Figure 3: Principal component analysis of network activity. A) Output after training a network to produce a sum of four sinusoids (red), and the approximation (brown) obtained using activity projected onto the 8 leading principal components. B) Projections of network activity onto the leading eight PC vectors. C) PCA eigenvalues for the network activity that generated the waveform in A. Only the largest 100 of 1000 eigenvalues are shown. D) Schematic showing the transition from control to learning phases of learning as a function of time and of PC eigenvalue. E) Evolution of the projections of $\mathbf{w}$ onto the two leading PC vectors during learning starting from five different initial conditions. These values converge to the same point on all trials. F) The same weight evolution but now including the projection onto PC vector 80 as a third dimension. The final values of this projection are different on each of the 5 runs, resulting in the vertical line at the center of the figure. Nevertheless, all of these networks generate the output in A.

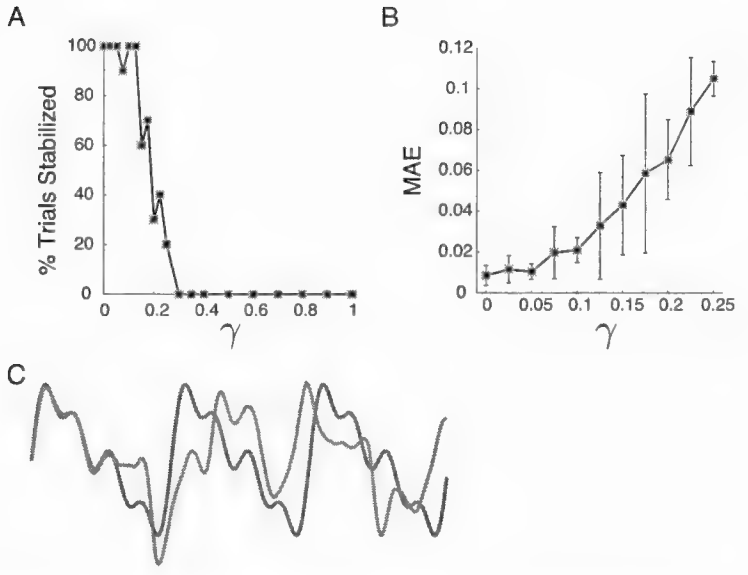

Figure 4: Comparison of different mixtures of FORCE ($\gamma = 0$ and echo-state ($\gamma = 1$) feedback. A) Percent of trials resulting in stable generation of the target function. B) Mean absolute error (MAE) between the output and target function after learning over the γ range where learning converged. C) Example run with output (red) and target function (green) for $\gamma = 1$. The trajectory is unstable.

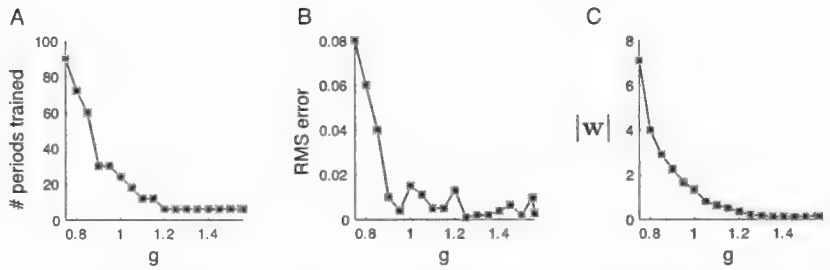

Figure 5: Chaos improves training performance. Networks with different g values (Methods) were trained to produce the output of figure 3A. Results are plotted against g in the range $0.75 < g < 1.56$, where learning converged. A) Number of cycles of the periodic target function required for training. B) The RMS error of the network output after training. C) The length of the readout weight vector $|\mathbf{w}|$ after training.

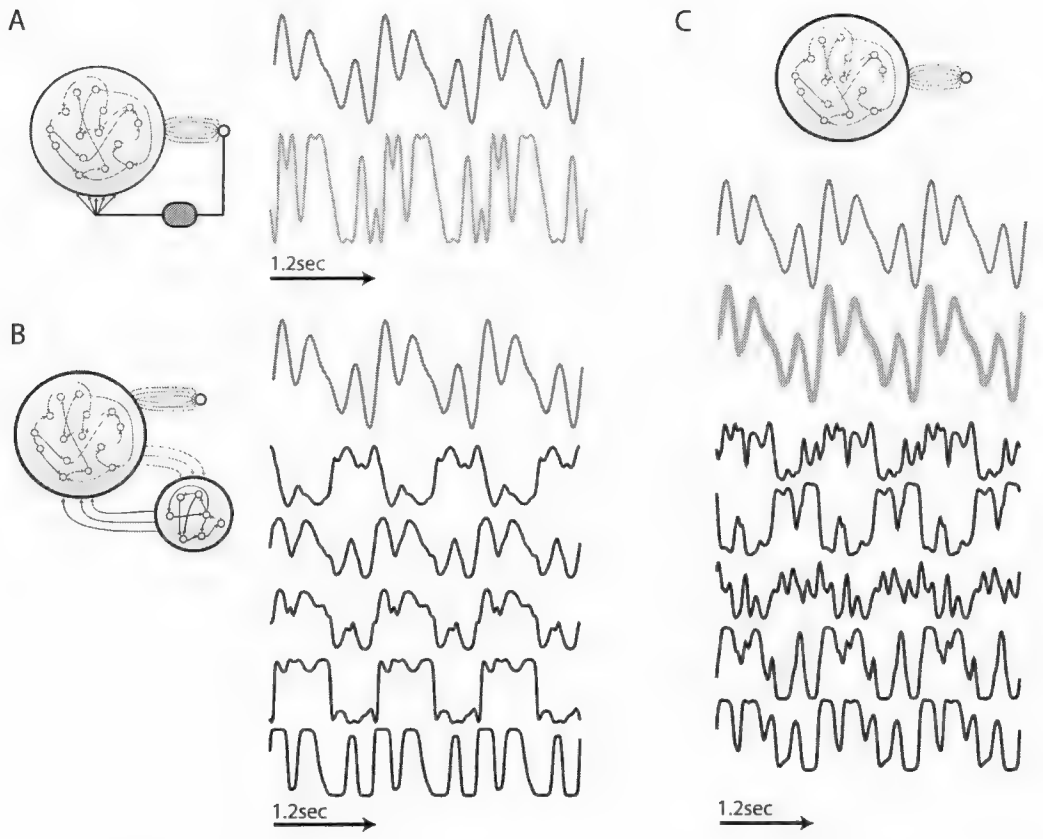

Figure 6: Feedback variants. A) Network trained to produce a periodic output (red trace) when its feedback (cyan trace) is $1.3\tanh(\sin(\pi z(t-100\text{ ms}))$, a delayed and distorted function of the output $z(t)$ (gray oval in circuit diagram). B) FORCE learning with a separate feedback network (circuit diagram). Output is the red trace, and blue traces show activity traces from 5 neurons within the feedback network. C) A network (circuit diagram) in which the internal synapses are trained to produce the output (red). Activities of 5 representative network neurons are in blue. The thick cyan traces are overlays of the component of the input to each of these 5 neurons induced by FORCE learning, $\sum_j \left(J_{ij}(t) - J_{ij}(0)\right) r_j(t)$ for $i = 1 \ldots 5$.

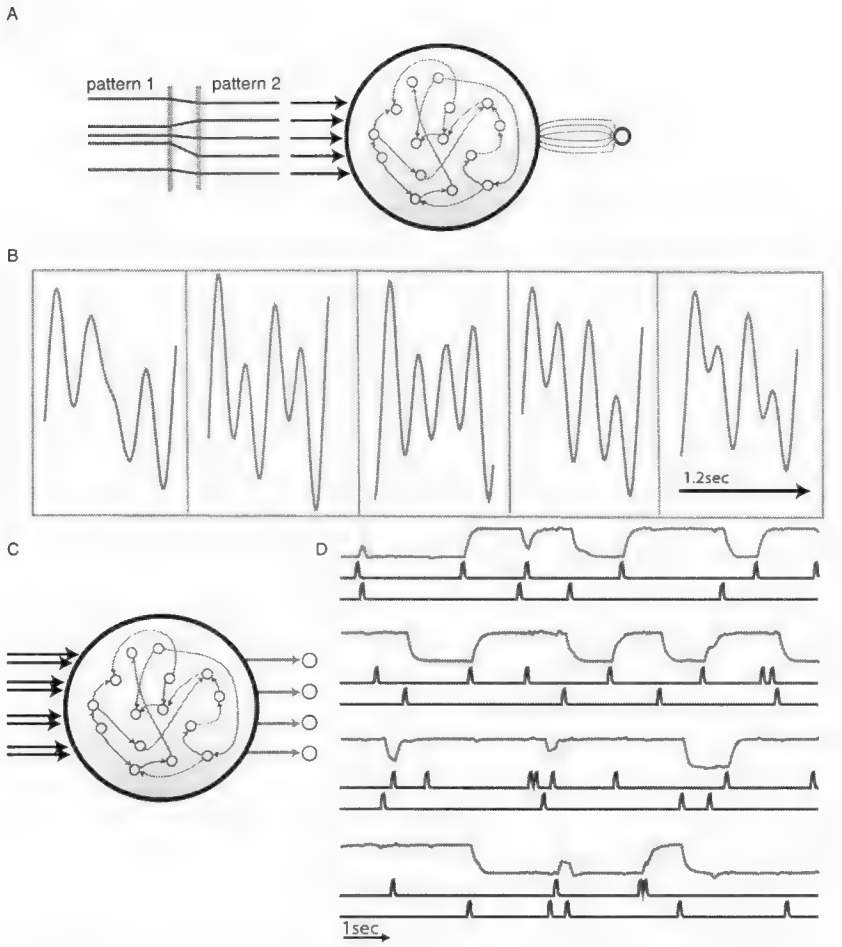

Figure 7: Multiple pattern generation and 4-Bit memory through learning in the generator network. A) Network with control inputs used to produce multiple output patterns (synapses and readout weights that are modifiable in red). B) Five outputs (1 cycle of each periodic function made from 3 sinusoids is shown) generated by a single network and selected by static control inputs. C) A network with 4 outputs and 8 inputs used to produce a 4-bit memory (modifiable synapses and readout weights in red). D) Red traces are the 4 outputs, with green traces showing their target values. Purple traces show the 8 inputs, divided into ON and OFF pairs associated with the output trace above them. The upper input in each pair turns the corresponding output on (sets it to +1). The lower input of each pair turns the output off (sets it to -1). After learning, the network has implemented a 4-bit memory, with each output responding only to its two inputs while ignoring the other inputs.

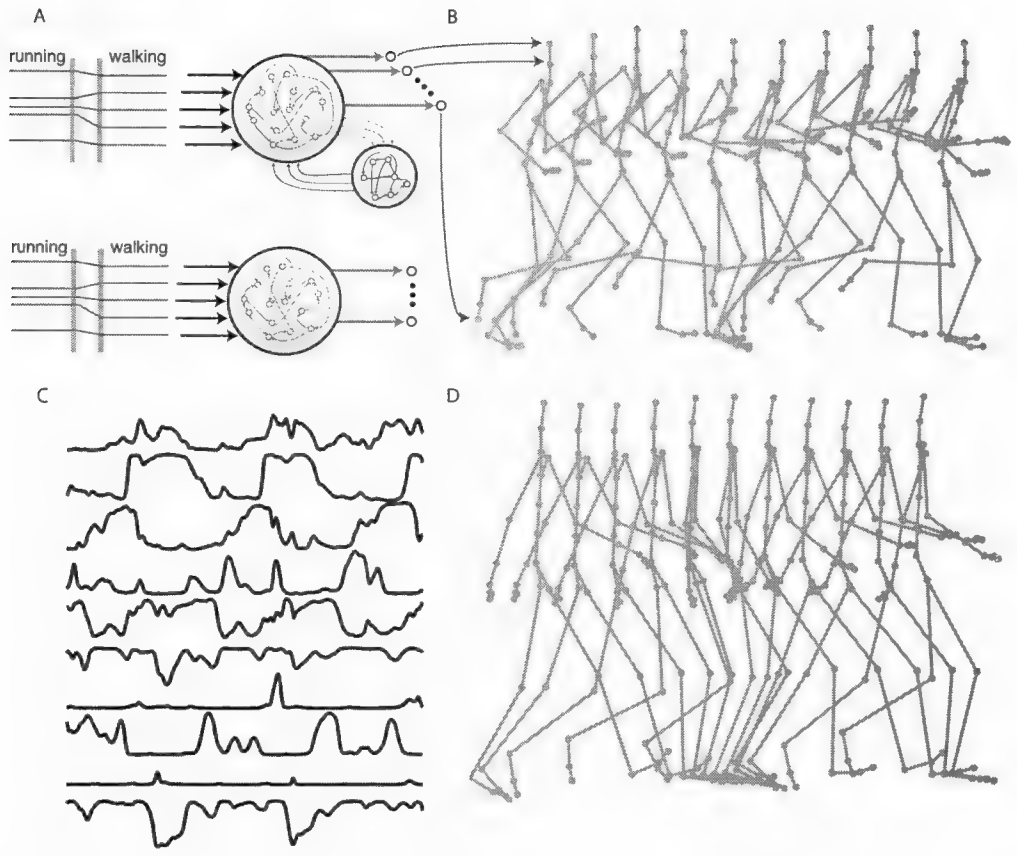

Figure 8: Networks that generate both running and walking human motions. A) Either of these two network architectures can be used to generate the running and walking motions (modifiable readout weights shown in red), but the upper network is shown. Constant inputs differentiate between running and walking (purple). Each of 95 joint angles is generated through time by one of the 95 readout units (curved arrows). B) The running motion generated after training. Cyan frames show early and magenta frames late movement phases. C) Ten sample network neuron activities during the walking motion. D) The walking motion, with colors as in B.

Annu. Rev. Neurosci. 2001. 24:139–66
Copyright © 2001 by Annual Reviews. All rights reserved

SYNAPTIC MODIFICATION BY CORRELATED ACTIVITY: Hebb's Postulate Revisited

Guo-qiang Bi[1] and Mu-ming Poo
Department of Molecular & Cell Biology, University of California at Berkeley, Berkeley, CA 94720-3200; email: mpoo@uclink4.berkeley.edu

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Key Words Hebbian synapse, LTP, LTD, spike timing, input specificity

■ **Abstract** Correlated spiking of pre- and postsynaptic neurons can result in strengthening or weakening of synapses, depending on the temporal order of spiking. Recent findings indicate that there are narrow and cell type–specific temporal windows for such synaptic modification and that the generally accepted input- (or synapse-) specific rule for modification appears not to be strictly adhered to. Spike timing–dependent modifications, together with selective spread of synaptic changes, provide a set of cellular mechanisms that are likely to be important for the development and functioning of neural networks.

> When an axon of cell A is near enough to excite cell B or repeatedly or consistently takes part in firing it, some growth or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.
>
> Donald Hebb (1949)

INTRODUCTION

Half a century since the publication of his famous treatise (Hebb 1949), Hebb's postulate of synaptic modification by correlated activity has become a cornerstone in our understanding of activity-dependent neural development and the cellular basis of learning and memory. This postulate was originally proposed by Hebb as a mechanism for the growth of "cell assembly," a hypothetical group of neurons that act briefly as a closed system after stimulation has ceased and that serve for the first stage of perception. Over the past several decades, Hebb's idea has been extended into various forms of correlation-based rules for synaptic modification and successfully used in many learning networks and in the analysis of activity-driven refinement of developing circuits (Stent 1973, Sejnowski & Tesauro 1989, Brown et al 1990, Fregnac & Bienenstock 1998, Sejnowski 1999). Here, we review

[1]Present address: Department of Neurobiology, University of Pittsburgh School of Medicine, Pittsburgh, PA 15261; e-mail: gqbi@pitt.edu

0147-006X/01/0301-0139$14.00

recent findings that shed new light on two aspects of Hebb's postulate: (*a*) the temporal specificity in the correlated activity required for the induction of synaptic modification—the importance of temporal order in the pre- and postsynaptic spiking; and (*b*) the spatial specificity in the induced synaptic changes—the notion that only the activated synapse becomes modified. Implications of these findings for the development and functioning of the nervous system are also addressed.

TEMPORAL SPECIFICITY IN ACTIVITY-INDUCED SYNAPTIC MODIFICATION

A central feature of Hebb's postulate is temporal specificity: The synaptic connection is strengthened only if cell A "takes part in firing" cell B, i.e. cell A fires before cell B. Such temporal specificity of activity-induced synaptic modification may be relevant for physiological functions, such as learning and memory, which are known to be temporally specific. Early studies have demonstrated a requirement for temporal contiguity in associative synaptic modification in rat hippocampus (Levy & Steward 1983) and *Aplysia* ganglia (Hawkins et al 1983, Walters & Byrne 1983). Recent experiments have revealed the importance of the temporal order of pre- and postsynaptic spiking for synaptic modification and have further defined the "critical windows" of spike timing, with precision on the order of milliseconds (Markram et al 1997, Magee & Johnston 1997, Bell et al 1997, Debanne et al 1998, Zhang et al 1998, Bi & Poo 1998, Egger et al 1999). The precise profile of critical windows appears to depend on the synapse type, and the underlying molecular mechanisms remain to be fully understood.

Temporal Requirements for Associative Synaptic Modification

Since the first description of long-term potentiation (LTP) in rabbit hippocampal formation (Bliss & Lømo 1973, Bliss & Gardner-Medwin 1973), similar synaptic potentiation has been found in many areas of the central and peripheral nervous systems of both vertebrates and invertebrates (Teyler & DiScenna 1987, Brown et al 1990, Bliss & Collingridge 1993, Buonomano & Merzenich 1998, Milner et al 1998, Malenka & Nicoll 1999). Certain forms of LTP in hippocampus have been shown to be cooperative and associative, properties relating directly to the theoretical construct of the Hebbian synapse. Cooperativity results from the existence of an intensity threshold for inducing LTP by tetanic stimulation (Bliss & Lømo 1973, McNaughton et al 1978), whereas associativity usually refers to the induction of LTP in a "weak" input when it is coactive with a "strong" convergent input (Levy & Steward 1979, Barrionuevo & Brown 1983). In a sense, cooperatively induced LTP is a special case of associative LTP: both originate from the requirement in the synchrony of inputs for postsynaptic activation (Teyler & DiScenna 1987).

Levy & Steward (1983) studied in more detail the temporal specificity in associative synaptic modification. By stimulating a weak and a strong input from the

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

entorhinal cortex to the denate gyrus of hippocampus, they found that associative induction of LTP did not require perfectly synchronous activation of the two pathways. Instead, the temporal order of the activation was crucial. LTP of the weak input could be induced when the strong input was activated concurrently with, or following the activation of, the weak input by as much as 20 ms. When the temporal order was reversed, long-term depression (LTD) was induced. This and other early studies (Kelso & Brown 1986, Gustafsson & Wigström 1986) have clearly indicated the existence of a stringent temporal specificity in the activity-induced synaptic modification.

In the marine mollusk *Aplysia*, temporally specific synaptic modifications have been shown to account for classical conditioning in several reflex systems (Carew et al 1981, Hawkins et al 1983, Walters & Byrne 1983). For example, in the siphon-withdrawal reflex, paired stimulation of the siphon sensory neuron [conditioned stimulus (CS)] and the tail nerve [unconditioned stimulus (US)] produces greater associative synaptic facilitation at sensorimotor connections when the CS precedes the US than if the temporal order is reversed (Hawkins et al 1983, Clark et al 1994). Such synaptic enhancement was considered non-Hebbian because the underlying cellular mechanism appeared to be purely presynaptic (Carew et al 1984). More recent studies, however, have shown postsynaptic involvement in associative synaptic facilitation of the same *Aplysia* system (Lin & Glanzman 1994, Murphy & Glanzman 1997). Temporally specific synaptic modification in invertebrates may thus be mediated by both presynaptic non-Hebbian and postsynaptic Hebbian mechanisms (Lechner & Byrne 1998).

The Role of Postsynaptic Spiking

Low-frequency presynaptic stimulation coupled with concurrent postsynaptic depolarization (the "pairing protocol") can induce LTP in the hippocampus and different cortical areas (Kelso et al 1986, Sastry et al 1986, Wigström et al 1986, Malenka & Nicoll 1999), demonstrating the importance of coincident activity in synaptic modification, as suggested by Hebb. However, this pairing protocol, with postsynaptic depolarization on the order of seconds to minutes, does not address the precise temporal specificity in the induction of LTP/LTD. Under natural conditions, postsynaptic neurons fire action potentials as their normal functional output, but whether these spikes are involved in inducing synaptic changes has been debated (for review, see Linden 1999). Active properties of dendrites (Regehr et al 1992, Lasser-Ross & Ross 1992, Johnston et al 1996) allow the spike initiated at the axon hillock to back-propagate into dendrites (Stuart & Sakmann 1994, Buzsáki et al 1996, Hoffman et al 1997). The back-propagating spike provides a precise signal capable of informing the synapses whether and precisely when the postsynaptic cell has fired, thus may play an active role in associative synaptic modification. The first definitive support for this notion came from recent studies in cortical and hippocampal slices (Markram et al 1997, Magee & Johnston 1997). Using dual whole-cell recording from two interconnected layer 5 pyramidal

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

neurons, Markram et al (1997) found that when the spikes were triggered 10 ms after the onset of excitatory postsynaptic potentials (EPSPs), LTP was induced by repetitive pairing of postsynaptic spiking (induced by current injection) with the EPSP, whereas when the temporal order of the spikes and EPSPs was reversed, LTD was induced. Magee & Johnston (1997) found that back- propagating spikes in CA1 pyramidal neurons of hippocampus, when coupled immediately after the onset of subthreshold EPSPs, evoked a significantly higher Ca^{2+} influx at the synaptic site, resulting in LTP. In both studies, neither EPSPs nor postsynaptic spiking alone was sufficient to induce any synaptic modification. Thus, back-propagating postsynaptic spikes indeed can function as an associative signal for synaptic modification, and relative timing of pre- and postsynaptic activity is critical. Similar results have also been obtained in slice cultures (Debanne et al 1998), in developing *Xenopus* retinotectal projections in vivo (Zhang et al 1998), and in cultured hippocampal neurons (Bi & Poo 1998).

Critical Windows of Spike Timing for Synaptic Modification

In cell culture, a critical window (Figure 1) for the induction of LTP/LTD has been characterized by systematically varying the spike timing (defined as the time interval between the onset of EPSP and postsynaptic spike) during repetitive correlated stimulation at a low frequency of 1 Hz (Bi & Poo 1998). The window for modification is about 40 ms in width and is temporally asymmetric. Postsynaptic spiking within about 20 ms after presynaptic activation (positive intervals) results in LTP, whereas that within about 20 ms before presynaptic activation (negative intervals) results in LTD. The transition from LTD to LTP occurs within a few milliseconds of change in spike timing. It is interesting that an essentially identical window has been found in developing retinotectal system of *Xenopus* tadpoles (Zhang et al 1998). At Schaffer collateral synapses in the CA1 region of rat hippocampal slices, LTP/LTD can also be induced by similar correlated pre- and postsynaptic activation at a low frequency (5 Hz), with a characteristic critical window of spike timing that differs slightly from that found in culture. An additional LTD window was observed at positive intervals of 15–20 ms (Nishiyama et al 2000).

In hippocampal cultures, GABAergic transmission can also be modified by correlated pre- and postsynaptic spiking. Such modification exhibits a symmetric spiking-timing window that lasts about 50 ms and is independent of *N*-methyl-D-aspartate receptors (NMDARs) (K Ganguly, M-m Poo, unpublished results). In the cerebellum-like structure of electric fish, synapses formed by parallel fiber onto Purkinje-like cells have a window opposite to that associated with the excitatory synapses described above: Spikes of positive intervals induce LTD and those of negative intervals induce LTP (Bell et al 1997). This opposite polarity in temporal specificity may be appropriate for the function this particular circuit, given that the Purkinje cell is an inhibitory projection neuron. In somatosensory cortex, excitatory neurons at different layers have markedly different windows. For synapses between layer 5 pyramidal neurons in somatosensory cortex, the

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

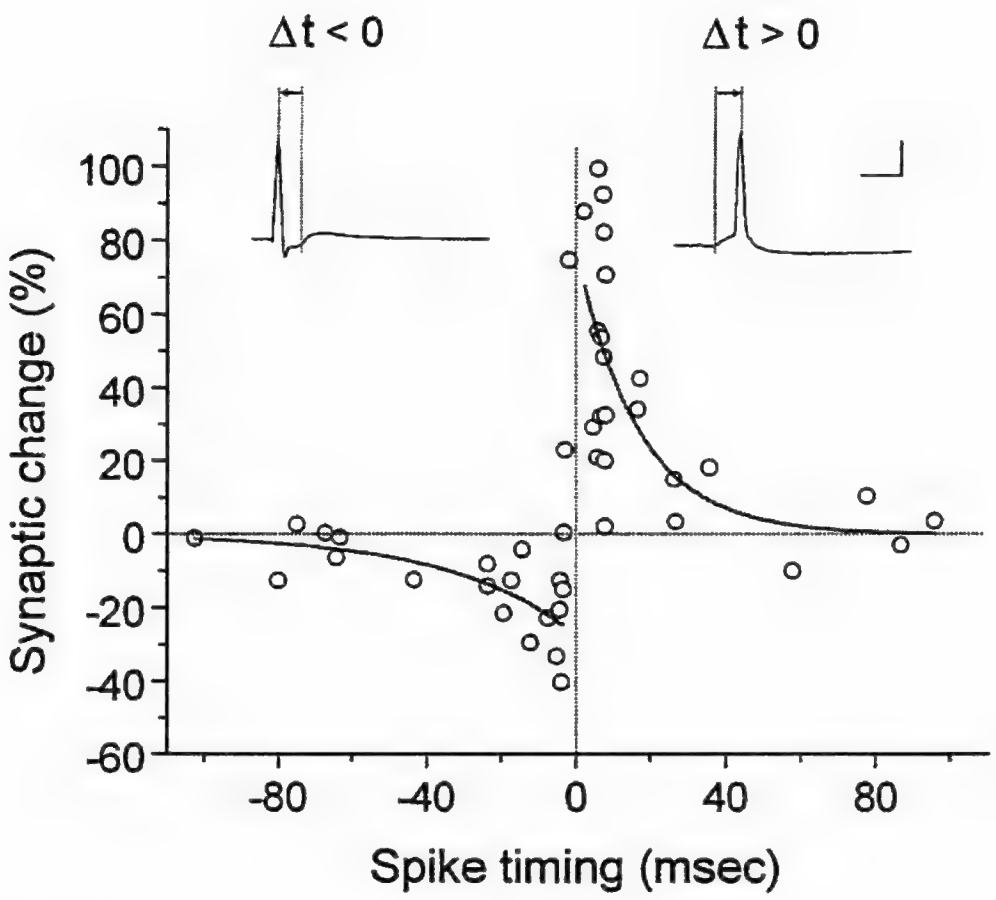

Figure 1 Critical window for synaptic modifications. Long-term potentiation (LTP)/long-term depression (LTD) were induced by correlated pre- and postsynaptic spiking at synapses between hippocampal glutamatergic neurons in culture. The percentage change in the excitatory postsynaptic current (EPSC) amplitude at 20–30 min after repetitive correlated spiking (pulses at 1 Hz) was plotted against spike timing, which is defined as the time interval (Δt) between the onset of the EPSP and the peak of the postsynaptic action potential during each pair of correlated spiking, as illustrated by the traces above. LTP and LTD windows are each fitted with an exponential function: $\Delta W = A \times e^{(-\Delta t/\tau)}$. For LTP and LTD, respectively, $A = 0.777$ and -0.273; $\tau = 16.8$ and -33.7 ms. Note that ΔW represents the total amount of change in synaptic strength after 60 pairs of correlated spiking. Assuming effective unitary change Δw (due to a single pair) has similar time constants as those of ΔW, i.e. $\Delta w = a \times e^{(\Delta t/\tau)}$, thus $(1 + A) = (1 + a)^{60}$, yielding $a = 0.0096$ and -0.0053 and the total areas under the unitary LTP and LTD curves $a \times \tau = 0.16$ and 0.18 for LTP and LTD, respectively. Scales: 50 mV, 10 ms. [Data from Bi & Poo (1998).]

window is presumably of similar profile as that in Figure 1 (Markram et al 1997). However, synapses between layer 4 spiny stellate neurons appear to have a symmetric depression window (Egger et al 1999). Various spike-timing windows for synaptic modification reflect the diversity of synaptic machinery and may serve for specific functions of information processing at different stages of neural pathways.

Mechanisms for the Detection of Coincidence and Temporal Order

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Previous studies on cellular mechanisms of LTP have addressed the role of the coincident pre- and postsynaptic activity in the induction of LTP (Bourne & Nicoll 1993). In some brain regions, including the CA1 area of hippocampus, synaptic activation of postsynaptic NMDAR, a subtype of glutamate receptor channel, concurrently with membrane depolarization underlies the coincidence requirement for LTP (for reviews, see Collingridge et al 1983, Bourne & Nicoll 1993). The depolarization-dependent removal of the Mg^{2+} block of the NMDAR (Mayer et al 1984, Nowak et al 1984, Jahr & Stevens 1987) allows the channel to serve as a molecular detector for the coincidence of presynaptic activation (glutamate release) and postsynaptic depolarization. The resultant Ca^{2+} influx through the NMDAR triggers downstream kinases that lead to synaptic potentiation. Induction of LTD also requires coincident pre- and postsynaptic activation and, in some cases, NMDAR activation, although lower-level postsynaptic depolarization is needed (Linden & Connor 1995, Bear & Abraham 1996, Goda & Stevens 1996). It is now generally believed that postsynaptic Ca^{2+} elevation is crucial for both LTP and LTD: Transient high-level Ca^{2+} elevation may lead to activation of certain protein kinases and LTP, whereas sustained low-level Ca^{2+} elevation may activate phosphatases and result in LTD (Lisman 1989, Malenka et al 1989, Malinow et al 1989, Silva et al 1992, Mulkey et al 1994, Yang et al 1999).

In principle, coincidence detection can occur at any converging point of signaling pathways where integration is nonlinear, as is the case for many biological processes. Thus, many components of signal transduction pathways may function as molecular coincidence detectors (Bourne & Nicoll 1993). In *Aplysia* sensorimotor synapse, presynaptic adenylyl cyclase that is modulated by G-protein–coupled serotinin receptors and Ca^{2+}/calmodulin detects the coincidence of Ca^{2+} elevation (due to sensory neuron activity, or CS) and serotonin release (from the facilitating interneuron due to US). The resultant rise in cAMP is at least partly responsible for the associative synaptic facilitation at this synapse (Abrams et al 1991, Yovell et al 1992, Lechner & Byrne 1998). In vertebrate brain, adenylyl cyclase is inhibited via G-protein by group II metobotropic receptors (mGluRs), which are known to be involved in several forms of LTP/LTD (Tzounopoulos et al 1998, Egger et al 1999, Yeckel et al 1999, Bortolotto et al 1999). Interaction between group II mGluR pathway and depolarization-induced Ca^{2+} signaling pathway is implicated in the coincidence detection at these forms of synaptic changes. Similarly, LTD in cerebellum and other forms of LTP/LTD in hippocampus require group I mGluRs that initiate Ca^{2+} release from intracellular stores via the activating

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

phospholipase C (Ito et al 1982, Linden & Connor 1995, Bear & Abraham 1996, Bortolotto et al 1999). A synergistic release of Ca^{2+} from inositol-1,4,5-triphosphate (IP3)-sensitive stores, evoked by synaptic activation of mGluRs, paired with back-propagating action potentials has been observed in the apical dendrites of CA1 pyramidal neurons (Nakamura et al 1999).

What cellular mechanisms are responsible for detecting the temporal order of pre- and postsynaptic spiking? Both LTP and LTD induced by correlated spiking depend on the activation of NMDARs (Markram et al 1997, Magee & Johnston 1997, Debanne et al 1998, Zhang et al 1998, Bi & Poo 1998), and LTD also requires functional L-type calcium channels (Bi & Poo 1998). These requirements are similar to LTP and LTD induced by the conventional tetanic stimulation or pairing protocols. For correlated spiking with positive intervals, spiking following synaptic inputs apparently helps to open NMDARs, allowing high-level Ca^{2+} influx, thus LTP. With negative intervals, spiking allows low-level Ca^{2+} influx through voltage-gated Ca^{2+} channels, which was followed immediately by additional low-level Ca^{2+} influx due to NMDARs, resulting in LTD. Indeed, imaging studies have shown that Ca^{2+} elevation was supralinear when presynaptic stimulation immediately preceded the spike and was sublinear when the temporal order was reversed (Koester & Sakmann 1998). Because the time constant of glutamate binding to NMDARs is much longer than 20 ms (Mayer et al 1984, Hollmann & Heinemann 1994), the narrow window observed in different systems must be accounted for by additional mechanisms, e.g. nonlinearity in the activation of downstream effectors (e.g. calmodulin), or EPSP-induced, dendritic A-type K^+ channel inactivation (time constant $\sim$10 ms) that permits synaptic invasion of back-propagating spikes (Hoffman et al 1997). Additional LTD at positive intervals found in hippocampal slices (Nishiyama et al 2000) may be due to differences in NMDAR properties, and K^+ channel activity, as well as the presence of local inhibitory inputs.

IMPLICATIONS OF TEMPORAL SPECIFICITY

The observation of stringent spike-timing windows highlights the potential role of individual spikes in synaptic modification (Linden 1999, Sejnowski 1999). The capability of synapses to detect precise timing of individual spikes is consonant with the idea that neural information may be encoded in spike timing (Perkel & Bullock 1968, Rieke et al 1997). Temporally asymmetric windows, such as that shown in Figure 1, appear to be the most predominant form in different systems (see Table 1). These windows provide a basis for the formulation of a spike-based, temporally asymmetric Hebbian learning rule (Gerstner et al 1996, Gerstner & Abbott 1997, Kempter et al 1999, Abbott & Song 1999, Rao & Sejnowski 2000, Mehta et al 2000, Senn et al 2000, Paulsen & Sejnowski 2000). Here, we focus on several properties of this new learning rule, namely the spike- rather than rate-based computation, self-normalization, and temporal asymmetry. Its functional implications for classical conditioning and temporal-to-spatial conversion of information in neural networks are also discussed.

TABLE 1 Spike timing (Δt) for the induction of LTP/LTD in different systems[a]

System	Δt LTD (ms)	Δt LTP (ms)	References
Hippocampal dentate gyrus	<0[b]	0 ~ +20	Levy & Steward 1983
Hippocampal slice CA1	None[c]	0 ~ +40 0 ~ +15	Gustafsson & Wigström 1986, Gustafsson et al 1987, Magee & Johnston 1997
	−30 ~ 10; +15 ~ +20	0 ~ +15	Nishiyama et al 2000
Cortical slice layer 5	−10[d]	+10[d]	Markram et al 1997
Cerebellum-like structure	0 ~ +50	−50 ~ 0	Bell et al 1997
Hippocampal slice culture	−100 ~ −15[d]	+15[d]	Debanne et al 1998
Retino-tectal projection	−20 ~ 0	0 ~ +20	Zhang et al 1998
Dissociated cell culture	−30 ~ 0	0 ~ +20	Bi & Poo 1998
Cortical slice layer 4	−10 ~ +10	None[e]	Egger et al 1999

[a]LTP, long-term potentiation; LTD, long-term depression.

[b]Heterosynaptic depression may contribute to LTD observed in this study.

[c]LTD was not observed or tested in the experiments.

[d]Single time points were tested for LTP/LTD in the experiments.

[e]No significant LTP was observed in the experiment.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Spike-Based, Temporally Asymmetric Learning Rule

A "Microscopic" Learning Rule One distinct feature of the spike-based learning rule is that, in the computation of synaptic modification, it emphasizes the precise timing of each individual spike rather than the average rate of a population of spikes over a defined interval. Thus, this rule can be considered "microscopic," as opposed to earlier statistical or "macroscopic" rate-based formulations, e.g. the covariance and the Bienenstock-Cooper-Munro (BCM) rules (Sejnowski 1977b, Bienenstock et al 1982), in both of which only the average spike rate is important. Using the spike-based rule and assuming particular firing characteristics of pre- and postsynaptic neurons, one can derive different statistical rules (Kempter et al 1999, Senn et al 2000) and readily incorporate them into more conventional rate-based models. An exciting application of spike-based learning rules is in a new type of neural network models, the "pulsed neural networks" (Maass & Bishop 1999), that use precise timing of individual spikes to encode information. Spike-timing or temporal coding may enable construction of more versatile and powerful networks because it provides larger coding capacity and easier handling of temporal information. However, whether the brain actually uses spike timing for coding information has been a long-standing issue in neuroscience (Perkel & Bullock 1968, Rieke et al 1997, Singer 1999). The critical question is, timewise, how precisely can a neuron (or a group of neurons) detect synaptic inputs and

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

generate as well as transmit spikes? The issues of precision and reliability of synaptic transmission, spike initiation, and propagation have been addressed by many theoretical and experimental studies (Calvin & Stevens 1968, Bullock 1970, Bekkers et al 1990, Mainen & Sejnowski 1995, Marsálek et al 1997, Stevens & Zador 1998, Diesmann et al 1999). It is interesting that theoretical analysis suggests that in multilayer networks, synchronous spiking activity may propagate from one ensemble to the next while maintaining high precision of spike timing (Marsálek et al 1997, Diesmann et al 1999). Such "mode-locked" spike propagation in the brain is consistent with the highly synchronous activity found in many areas of the brain (Abeles et al 1993, deCharms & Merzenich 1996, Riehle et al 1997, Roelfsema et al 1997, Dan et al 1998, Singer 1999). The existence of narrow windows of spike timing for synaptic modification indicates that synapses can indeed "read" information coded in the timing of pre- and postsynaptic spikes, with a precision on the order of a few milliseconds. Therefore, spike timing–based learning rules and the pulsed network models may provide powerful tools as well as insights into computational principles of the nervous system.

A "Self-Normalizing" Rule Hebb's original postulate described only activity-dependent strengthening of a synapse. However, a complete learning rule must include both strengthening (LTP) and weakening (LTD) of synapses in order to avoid saturation, and weakening must be as specific as strengthening (rather than nonspecific decay) for serving long-term memories (Sejnowski 1977a, Stevens 1996). The new asymmetric Hebbian rule provides an immediate mechanism for self-normalization of synaptic weights and the output firing rate of a network (Kempter et al 1999, Abbott & Song 1999, Senn et al 2000). Under normal conditions, such a balance is achieved if, in the spike-timing window of unitary modification (the effect of a single pairing event), the total integral of the LTD curve is equal to or slightly larger than that of the LTP curve. The data shown in Figure 1 indeed largely satisfy this requirement of self-normalization. In some synapses, the spike-timing windows for correlated spiking-induced modification appear not self-normalized (Egger et al 1999). It is possible that there is a low-level (undetectable) timing-independent LTD/LTP or activity-independent LTD/LTP that serves for nonspecific normalization of these synapses.

A "Temporally Asymmetric" Rule Perhaps the most attractive feature of the new learning rule is its sharp asymmetry in the window of spike timing. Association between two events of a specific temporal order can be established by asymmetric synaptic modification. This is fundamentally different from most formulations of Hebbian rules, in which only coincidence of activity (or associated between two concomitant events) was considered (Brown et al 1990). Hebb's original statement was remarkably accurate in implying the role of temporal order: A synaptic input is strengthened only when it "takes part in firing" the postsynaptic cell. Thus, a "true" Hebbian synapse behaves as a causality detector rather than a coincidence detector. Networks of neurons with such synapses can learn and predict sequences (Minai & Levy 1993, Abbott & Blum 1996, Roberts 1999) and form navigational

maps that have been observed in hippocampus (Mehta et al 1997, Gerstner & Abbott 1997, Mehta et al 2000). More generally, the predictive nature of the temporally asymmetric learning rule is directly related to the implementation of classical conditioning (see below).

Classical Conditioning

A hallmark of classical conditioning is the requirement of temporal contiguity: the conditioned stimulus (CS) must precede the unconditioned stimulus (US) during conditioning in order to elicit conditioned response. This ability to detect and remember causal relationships between events, to predict the about-to-happen, and to respond in advance are obviously advantageous to all animal species. Classical conditioning has been modeled successfully in the past by temporal-difference learning algorithm in reinforcement learning (Sutton & Barto 1981, Tesauro 1986, Montague & Sejnowski 1994, Moore et al 1998). It has been shown recently that the asymmetric Hebbian rule is mathematically equivalent to the temporal-difference learning algorithm (Roberts 1999, Rao & Sejnowski 2000). Therefore, the asymmetric synaptic modification by correlated spiking provides a natural cellular mechanism of temporal-difference learning at the synaptic level.

Classical conditioning normally operates over a timescale on the order of seconds. In *Aplysia* siphon-withdrawal reflex, temporally specific synaptic enhancement at sensorimotor connections can be induced when stimulation of the tail (US) follows stimulation of the siphon sensory neuron (CS) within seconds (Hawkins et al 1983, Walters & Byrne 1983, Clark et al 1994). In this case, cellular process with slow kinetics at the presynaptic terminal—the interaction between Ca^{2+} and adenylyl cyclase—provides a detection mechanism for a temporal order of sensory events (Yovell & Abrams 1992). In vertebrate brain, processes such as synergistic release of postsynaptic Ca^{2+} stores by mGluR activation paired with back-propagating spikes (Nakamura et al 1999) may also provide temporally specific detection with a longer timescale than that mediated by NMDARs. Does correlated-spiking–induced synaptic modification found for hippocampal and cortical synapses, with the spike-timing window lasting tens of milliseconds, play a role in more elaborated forms of classical conditioning in vertebrates known to have temporal windows that last for seconds (Squire 1987, Thompson & Krupa 1994)? Many different areas are likely to be involved in classical conditioning within the vertebrate brain, with polysynaptic and recurrent circuitry that can effectively extend the timescale of computation. For example, polysynaptic pathways may introduce long transmission delays so that a "remote" synapse may detect and become modified by two temporally distant stimuli that arrive at the site of modification within a narrow window of spike timing (Bi & Poo 1999).

Temporal-to-Spatial Conversion: Delay-Line Mechanism

An essential process in learning and memory is the transformation of temporal information into spatially distributed information in the brain. Temporally varying

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

sensory information, e.g. that associated with vision and audition, are processed and stored by the brain with remarkable precision (Carr 1993, Singer & Gray 1995). The essential temporal information may be coded in the timing of transmitted spike trains (Rieke et al 1997, Singer 1999). In addition, inputs that are not explicitly temporal, e.g. odorant stimuli, may also be coded in spike timing (Hopfield 1995, Laurent 1997). Relative timing of spikes among neurons within functional ensembles may be used directly to represent internal states of functioning circuits (Riehle et al 1997). For synaptic modification to be a mechanism for long-term storage of memory, such information coded in spike timing must be converted into and stored as spatially distributed synaptic modifications in the brain. This requires that each synapse be sensitive to spike timing, as discussed above, and that a mechanism "assign" the temporal information specifically to different synaptic locations. One strategy for accomplishing this conversion relies on intrinsic synaptic properties (Buonomano & Merzenich 1995), such as short-term plasticity and slow inhibitory currents that provide differential neuronal sensitivity to different spikes in an input train (Buonomano et al 1997). Another strategy is to take advantage of network architectures of "delay-lines" (Jeffress 1948, Braitenberg 1967, Tank & Hopfield 1987) that have been demonstrated in several systems where the delay is due to axonal conduction of action potentials (Carr 1993). Recently, we have shown that in networks of cultured neurons, polysynaptic transmission pathways may form functional delay-lines (Bi & Poo 1999). Repetitive paired-pulse stimulation of one neuron results in correlated spiking and LTP/LTD at remote synaptic sites, when the interpulse interval matches the delay difference between convergent pathways leading to the remote sites (Bi & Poo 1999). Because neurons in such cultures tend to hyper-innervate each other, these polysynaptic pathways in culture may be analogous to chains of synchronously firing cell ensembles in vivo that have been proposed for cortical processing (Abeles 1991). With such delay-lines and spike-timing dependent LTP/LTD, information coded in the spike train may be stored at selective connections between cell ensembles in the form of long-term synaptic modification.

Activity-Driven Refinement of Developing Networks

Since the early works of Hubel and Wiesel on the effects of monocular deprivation (Wiesel & Hubel 1963) and artificial strabismus (Hubel & Wiesel 1965) on the developing visual system, it has been generally recognized that neuronal activity is critical for refining the connections in the brain (for reviews, see Katz & Shatz 1996, Constantine-Paton et al 1990). There is also evidence that activity may provide an instructive (rather than permissive) role. For example, ocular dominance columns in the cortex were abolished if artificially synchronized activity was imposed to both eyes, whereas imposing asychonous activity between the eyes of the same frequency was not effective (Stryker & Strickland 1984). Development of orientation selectivity and responsiveness of visual cortical neurons can also be reduced by artificial stimulation of optic nerves (Weliky & Katz 1997). Although

patterned visual experience may not be necessary for the initial development of either ocular dominance or orientation maps (Horton & Hocking 1996, Crair et al 1998, Crowley & Katz 1999), the endogenous patterns of spontaneous activity (Wong et al 1995, Penn et al 1998), in the retina as well as in higher visual areas, could still play a role.

In addressing the cellular mechanism by which asynchronous activity between the two eyes drives the formation of ocular dominance columns, Stent (1973) proposed an extension of Hebb's postulate: Synchronous activity strengthens the synapse, whereas asynchronous activity weakens the synapse. This simple rule of "neurons that fire together wire together" can explain activity-dependent competitive interactions between converging inputs on the same postsynaptic neurons. Synchronous spiking is more likely to occur for inputs from the same eye; thus, there is eye-specific segregation of inputs in lateral geniculate nucleus and in the visual cortex. Correlation of spiking is stronger for inputs from neighboring neurons in the same retina; thus, a crude retinotopic map (presumably formed initially by chemical cues) can be refined into a more fine-grained map (Constantine-Paton et al 1990). Simple synchrony- or correlation-based rules have also been successfully used in modeling activity-driven development of simple receptive fields and the ordered arrangement of ocular dominance and orientation columns (Miller et al 1989, Miller 1994), as well as ocularly matched orientation and ocular dominance maps (Erwin & Miller 1998).

The finding of spike-timing–dependent modifications raises several questions. Is a simple correlation-based Hebbian rule a realistic or complete description of activity-driven processes? Are there activity-dependent processes other than the modification of synaptic strength that are correlation based but spike-timing independent? Does introduction of spike-timing dependency into correlation-based rule offer new features for the more complex aspects of circuit developments? It is interesting to note that spike-timing–dependent rules have been used successfully in modeling the development of fine temporal discrimination in the auditory system (Gerstner et al 1996) and direction-sensitive receptive field organization of mammalian visual cortical neurons (Rao & Sejnowski 2000).

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

SPATIAL SPECIFICITY IN ACTIVITY-INDUCED SYNAPTIC MODIFICATION

Although not explicitly stated in Hebb's postulate, it is generally assumed that Hebbian synaptic modifications are synapse (or input) specific—only synapses experiencing correlated activity become modified. Input specificity was indeed observed in many studies of LTP/LTD (Andersen et al 1977, Lynch et al 1977, Dudek & Bear 1992, Mulkey & Malenka 1992, Nicoll & Malenka 1997). Thus, Hebbian rule is regarded as a "local" rule, in which individual synapses are independent of one another. Although the simplicity of this local rule has been

TABLE 2 Spread of LTP/LTD observed in different systems[a]

Induction	Spread[b]	Systems (References)
LTP	Presynaptic lateral (diverging outputs)	Hippocampal slice CA1 (Schuman & Madison 1994)
		Hippocampal slice culture (Bonhoeffer et al 1989)
		Visual cortex slice (Kossel et al 1990)
		Hippocampal culture (Tao et al 2000)
	Postsynaptic lateral (converging inputs)	Hippocampal slice CA3 (Bradler & Barrioneuvo 1989)
		Hippocampal slice CA1 (Muller et al 1995)
		Hippocampal slice culture (Engert & Bonhoeffer 1997)
		Hippocampal aspinous interneurons (Cowan et al 1998)
	Back propagation	Hippocampal culture (Tao et al 2000)
LTD	Presynaptic lateral	Hippocampal culture (Fitzsimonds et al 1997)
	Postsynaptic lateral	Hippocampal slice CA1 (Muller et al 1995, Staubli & Ji 1996, Nishiyama et al 2000)
		Hippocampal aspinous interneurons (Cowan et al 1998)
		Hippocampal interneurons (McMahon & Kauer 1997)
		Hippocampal culture (Fitzsimonds et al 1997)
	Back propagation	Hippocampal culture (Fitzsimonds et al 1997)

[a]LTP, long-term potentiation; LTD, long-term depression.

[b]See text and Figure 2 for definitions.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

beneficial for many neural network models (Brown et al 1990, Rolls & Treves 1998), there is now increasing evidence that activity-induced synaptic modification may be accompanied by changes in some other synapses within a neural network (see Table 2).

Heterosynaptic LTD

Early studies of LTP in the CA1 region of hippocampus have shown a depression of unstimulated commissural-CA1 synapses following LTP-inducing tetanic stimulation of Schaffer collaterals (Lynch et al 1977, Dunwiddie & Lynch 1978). This phenomenon, termed heterosynaptic LTD, has since been observed in the cortex (Tsumoto & Suda 1979, Hirsch et al 1992), dentate gyrus (Levy & Steward 1979, Abraham & Goddard 1983), and the CA3 area (Bradler & Barrioneuvo 1989) of hippocampus, as well as at embryonic neuromuscular junctions in culture (Lo & Poo 1991). The conditions for the induction and expression of heterosynaptic LTD vary between different systems, but appear to be closely related to that for homosynaptic LTD (for reviews, see Linden & Connor 1995, Bear & Abraham 1996). A moderate rise in postsynaptic Ca^{2+}, through influx mediated by NMDARs and other Ca^{2+} channels as well as through Ca^{2+} release from internal stores, has been implicated in the induction mechanism, although the signaling pathways downstream of Ca^{2+} elevation and the mechanism of expression remain

largely unknown. Heterosynaptic LTD induced by the induction of LTP represents perhaps a special case of more general non-local effects of synaptic modification (see below).

Spread of LTP/LTD in Slices

Although heterosynaptic LTD indicates the existence of interaction between distant synapses, it is not considered a violation of input specificity because changes at nonstimulated sites are opposite in polarity from those at induction sites. However, exceptions to this rule have been observed since the early days of LTP studies (for review, see Teyler & DiScenna 1987). For example, in area CA3 of hippocampus, mossy fiber tetanization leads to heterosynaptic LTP, as reflected in an increased field potential evoked by other converging but independent inputs, in addition to the conventional homosynaptic LTP (Yamamoto & Chujo 1978, Misgeld et al 1979). This heterosynaptic LTP may be caused by an increase in the polysynaptic component of CA3 field potentials, resulting from enhanced transmission among CA3 neurons (Higashima & Yamamoto 1985). It is interesting that Bradler & Barrioneuvo (1989) found that among three inputs to the CA3 area, tetanus-induced LTP in the mossy fiber input was accompanied by heterosynaptic LTP of Schaffer collateral and fimbrial responses, whereas LTP in the Schaffer pathway was associated with the LTP of fimbrial responses and LTD of mossy fiber responses. In contrast, LTP induced in the fimbrial response was input specific.

More recent works have revealed different forms of "breakdown" of synapse specificity of LTP in the CA1 region of hippocampus. In slice cultures, Bonhoeffer et al (1989) found that LTP induced by pairing presynaptic stimulation with postsynaptic depolarization is accompanied by potentiation at synapses made by the stimulated fibers on neighboring postsynaptic cells. Similar "spread" of potentiation was also observed in the visual cortex, where the spread of potentiation was shown to occur only for synapses made by the stimulated, but not unstimulated, inputs (Kossel et al 1990). A closer examination using local perfusion indicated that, in hippocampal slice cultures, LTP can spread to other inputs on the same postsynaptic cell as long as the unstimulated synapses are within $\sim$70 μm of the site of LTP induction (Engert & Bonhoeffer 1997). Also observed at excitatory synapses was postsynaptic spread of LTP onto aspinous dendrites of interneurons whose somata were located in the CA1 pyramidal cell layer (Cowan et al 1998). It is interesting that pairing-induced LTP in CA1 pyramidal cells of acute hippocampal slices also spread to synapses made by the same set of afferent fibers onto neighboring postsynaptic cells (Schuman & Madison 1994). Taken together, these studies revealed two forms of spread of LTP. In one form, synaptic sites of induction and spread are apparently made by the same (group of) presynaptic axons on different postsynaptic cells (Bonhoeffer et al 1989, Kossel et al 1990, Schuman & Madison 1994). In the other form, the synapses of induction and spread are on the same postsynaptic cell but made by different presynaptic axons (Engert & Bonhoeffer 1997, Cowan et al 1998). For convenience, we

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

define these two forms presynaptic spread and postsynaptic spread, respectively (Figure 2). Note that presynaptic spread can only be defined unequivocally when the stimulation is applied to a single presynaptic neuron. Thus, the cases using extracellular stimulation of afferent fibers (Bonhoeffer et al 1989, Kossel et al 1990, Schuman & Madison 1994) can only be regarded as consistent with the definition of presynaptic spread.

Several studies of conventional homosynaptic LTD have also suggested a breakdown of synapse specificity under different conditions. In CA1, homosynaptically induced LTD was found to be associated with a heterosynaptic reversal of the LTP previously induced in a separate pathway, and the effect was absent for naïve inputs and dependent on both NMDAR and calcineurin (Muller et al 1995). Another study, however, suggested that LTD at one of two independent Schaffer-collateral/commissural inputs may spread to the other (presumably naïve) input when the level of synaptic activation during low-frequency stimulation was strong enough to induce postsynaptic spiking (Staubli & Ji 1996). Spread of a different form of LTD has been observed in excitatory synapses onto GABAergic interneurons in the CA1 area of hippocampus: Tetanic stimulation of the input fiber resulted in LTD that spread to neighboring excitatory synapses onto the same postsynaptic interneurons (McMahon & Kauer 1997). In all these cases, only postsynaptic spread has been examined. These results differ from conventional heterosynaptic LTD in that the change at the induction pathway is also LTD rather than LTP.

The issue of input specificity of LTP/LTD induced by correlated pre- and postsynaptic spiking has been examined recently for Schaffer collateral-CA1 synapses in hippocampal slices (Nishiyama et al 2000). It was found that correlated activity of positive intervals induces input-specific homosynaptic LTP, whereas that of negative intervals induces LTD that spreads to heterosynaptic unstimulated inputs. Reduction of postsynaptic Ca^{2+} influx by partial blockade of NMDA receptors resulted in a conversion of LTP to LTD and a "breakdown" of input specificity, with LTD appearing at heterosynaptic inputs. The induction of LTD at homo- and heterosynaptic sites requires functional ryanodine receptors and IP_3 receptors (IP_3R), respectively. Functional blockade or genetic deletion of type 1 IP_3R led to a conversion of LTD to LTP and elimination of heterosynaptic LTD, whereas blocking ryanodine receptors eliminates only homosynaptic LTD. These results indicate that both the polarity and input specificity of activity-induced synaptic modifications are tightly regulated by postsynaptic Ca^{2+}, derived from Ca^{2+} influx and differential release of Ca^{2+} from internal stores. Therefore, input specificity should be viewed not as an intrinsic property associated with LTP/LTD but as a dynamic variable linked to the spread of dendritic Ca^{2+} elevation.

Selective Spread of LTP/LTD in Cell Cultures

Recent studies on synaptic plasticity in cell cultures have revealed extensive but selective spread of both LTP and LTD from the site of induction to other synapses

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

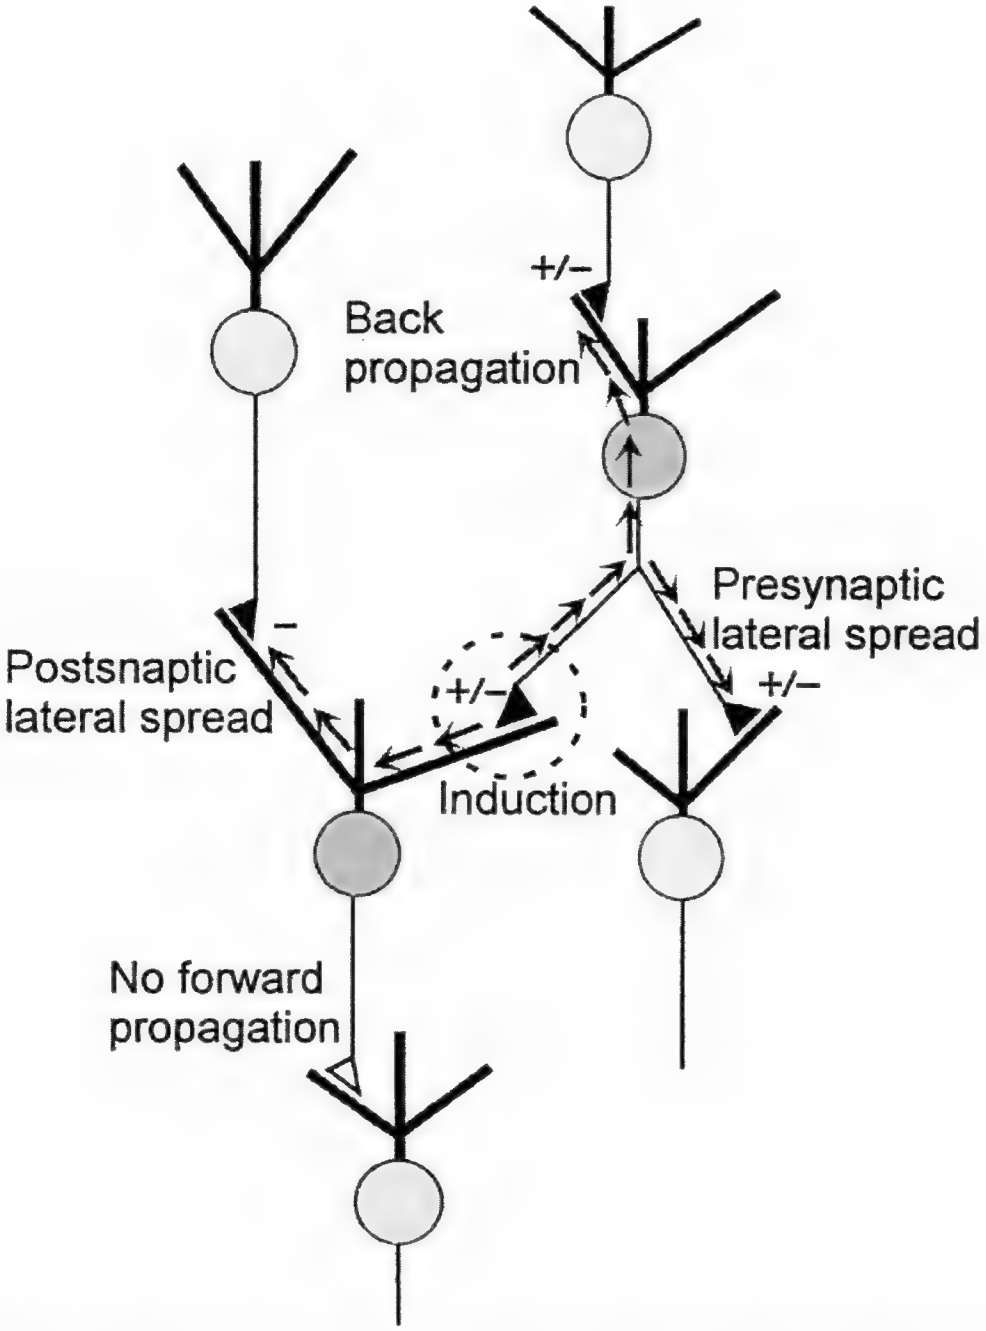

Figure 2 Different types of spread of long-term potentiation (LTP) and long-term depression (LTD). Induction of LTP/LTD at a synapse between glutamatergic neurons (black triangle in dashed circle) leads to the spread of LTP/LTD (marked as +/−) to other synapses in the network. Arrows indicate direction of signaling underlying the spread of LTP/LTD to divergent outputs of the presynaptic neuron (presynaptic lateral spread) and synapses on the dendrites of presynaptic neuron (back propagation), as well as the spread of LTD to convergent inputs on the postsynaptic neuron (postsynaptic lateral spread). Spread of LTP/LTD to the output of postsynaptic neurons (forward propagation) has not been observed.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

in a network. In *Xenopus* nerve-muscle cultures, LTD induced by postsynaptic elevation of Ca^{2+} can spread to synapses made by the same presynaptic neuron on other myocytes, apparently by signaling within the cytoplasm of the presynaptic neuron (Cash et al 1996). In small networks of cultured hippocampal neurons, LTD induced at synapses between two glutamatergic neurons can spread to other synapses made by divergent outputs of the same presynaptic neuron (presynaptic lateral propagation) and to synapses made by other convergent inputs on the same postsynaptic cell (postsynaptic lateral propagation). Furthermore, LTD can also spread in a retrograde direction to depress synapses on the dendrite of the presynaptic neuron (back propagation) but not in anterograde direction (forward propagation) to depress output synapses of the postsynaptic neuron (Fitzsimonds et al 1997) (see Figure 2). It is interesting that in the same culture system, LTP induced at synapses between two glutamatergic neurons by correlated spiking exhibits a selective spread only to those synapses associated with the presynaptic neuron (i.e. both lateral and back propagation), not to those associated with the postsynaptic neuron (no lateral or forward propagation) (Tao et al 2000). In addition, the presynaptic lateral propagation of LTP is target cell specific: Divergent outputs to GABAergic postsynaptic neurons are not modified. This specificity is similar to the target specificity found in the induction of LTP in both cultures and slices (McMahon & Kauer 1997, Maccaferri et al 1998, Bi & Poo 1998, Reyes et al 1998, Schinder et al 2000). Finally, neither LTP nor LTD was found to spread further up- or downstream beyond the immediate neighboring synapses (Fitzsimonds et al 1997, Tao et al 2000). Although findings in cell culture may represent an exaggeration of phenomena occurring in vivo, the simplicity of the culture system enables investigation of cellular processes and network principles that may be present to varying degrees in more complex neural systems.

Development of Input Specificity

Recent studies of activity-induced synaptic modification in the developing *Xenopus* retinotectal system have suggested a cellular basis for the input specificity associated with LTP (H-z Tao, L Zhang, F Engert, M-m Poo, unpublished results). The induction of LTP at retinotectal synapses by theta burst stimulation of the retinal ganglion cells was not input-specific: LTP induced in one retinal ganglion cell pathway spreads to other converging retinal ganglion cell inputs on the same tectal neuron. Inputs on other adjacent tectal neurons were not affected, which suggests that the spread of potentiation was due to signaling within the postsynaptic cytoplasm. As the animal matures, LTP of retinotectal synapses induced by the same theta burst activity becomes input specific, a change that correlates with increased complexity of the dendritic arbor of tectal neurons and more restricted distribution of dendritic Ca^{2+} elevation evoked by retinal inputs. In contrast to that found for theta burst activity, LTP induced by low-frequency correlated pre- and postsynaptic spiking was input specific throughout the development. These results showed that a "break-down" of input specificity of

activity-induced LTP can occur during in vivo development. The extent of the spread of synaptic potentiation, which correlates with spatial distribution of dendritic Ca^{2+} elevation, depends on the developmental stage as well as the pattern of synaptic activity.

IMPLICATIONS OF SPATIAL SPECIFICITY (OR LACK OF IT)

Long-Range Signaling in and Across Neurons

The spread of LTP/LTD and heterosynaptic LTD indicates that there are interactions among different synapses by long-range signaling either within the neuronal cytoplasm or across the neurons through the extracellular space. It has been suggested that membrane-permeable diffusible factors are responsible for retrograde signaling during the induction of LTP and LTD (for review, see Bliss & Collingridge 1993, Fitzsimonds & Poo 1998). These diffusible factors may affect other synapses close to the sites of induction. Schuman & Madison (1994) have suggested that nitric oxide may mediate the spread of LTP in hippocampus because inhibition of nitric oxide synthase in the postsynaptic cell blocked induction of LTP as well as its spread. However, a different signal triggered by the retrograde signal associated with LTP (e.g. nitric oxide) could be responsible for the spread of potentiation, and the signal could be confined intracellularly.

The extensive but highly selective spread of LTP/LTD found in culture argues for the existence of a rapid long-range intracellular signaling within the cytoplasm, with a speed of at least a few micrometers per second. This requires cytoskeleton-based axonal transport or regenerative waves as a means of signaling. Induction of LTP/LTD may generate Ca^{2+} waves (Berridge 1998) in either pre- or postsynaptic neuron and, through its diverse downstream effectors, results in either heterosynaptic LTD or spread of LTP/LTD. At Schaffer collateral-CA1 synpases, LTD induced by correlated pre- and postsynaptic activation (of negative intervals) spreads to unstimulated inputs to the same CA1 neuron (Nishiyama et al 2000). This spread was abolished by postsynaptic loading of a specific function-blocking antibody of type 1 IP_3R and was absent in hippocampal slices obtained from IP_3R-deficient mice, consistent with the notion that an IP_3R-dependent wave may be responsible for postsynaptic spread of LTD. Another second-messenger cAMP is known to be involved in different forms of synaptic plasticity and can diffuse rapidly in neuronal cytoplasm (Hempel et al 1996). Local elevation of cAMP/PKA activity in a developing neuron has been shown to exert long-range actions on distant parts of the neuron (Zheng et al 1994). In *Aplysia*, restricted application of serotonin to the cell bodies of sensory neurons can cause long-term enhancement at synaptic connections of the same sensory neuron to distant motoneurons

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

(Clark & Kandel 1993, Emptage & Carew 1993). Thus, cAMP or its downstream effectors can serve for cytoplasmic signaling involved in the spread of synaptic modification.

Long-range cytoplasmic signaling may also involve motor protein-based transport of vesicular membranes (Vale et al 1985, Vallee & Bloom 1991, Kuznetsov et al 1992, Bi et al 1997). Fast axonal transport of vesicles carrying internalized retrograde factors (e.g. neurotrophins) or vesicle-associated cytoplasmic effectors triggered by the induction of LTP/LTD may spread the signal for potentiation/depression within the cytoplasm. Restricted application of glutamate to somata of identified presynaptic neurons led to LTP of its output synapses, and this LTP can be blocked by pretreatment of microtubule-depolymerizing reagent colchicine (Lux & Veselovsky 1994). However, depolymerization of microtubule may also disrupt the normal structure of endoplasmic reticulum that supports the propagation of Ca^{2+} waves. Axotomy of postganglionic fibers leads to regression of synaptic inputs from preganglionic axon, a phenomenon that can be simulated by local chochicine block of the postganglionic fiber and can be prevented by supplying an exogenous neurotrophin nerve growth factor (Purves 1975). In hippocampal cell cultures, brain-derived neurotrophic factor induces synaptic potentiation by elevating transmitter secretion in a target cell–specific manner (Schinder et al 2000), i.e. only for those terminals innervating glutamatergic but not GABAergic neurons. This specificity has also been seen in the induction and spread of LTP in the same cultures (Bi & Poo 1998, Tao et al 2000). Thus, brain-derived neurotrophic factor may also be a good candidate for the cytoplasmic signal for presynaptic spread of LTP.

Non-Local Rules in Neural Networks

Back-Propagation Although synapse specificity has been a "central dogma" in the concept of synaptic modification, a popular network model of supervised learning adopted a non-local rule, with which signals for synaptic modification can propagate backward through connections in multilayer networks (Rumelhart et al 1986a,b). This "back-propagation" algorithm, which uses the "error" of the output of a neuron to proportionally adjust the synaptic strength of its inputs, has been extremely powerful in training simple artificial networks to perform a wide range of tasks. The back-propagation algorithm has been considered nonbiological because it directly violates the rule of synapse specificity and is inconsistent with the known direction of signaling in neurons, i.e. from the input to the output (cf Rolls & Treves 1998). It is interesting that back propagation of LTP/LTD observed in cell cultures appears to fit qualitatively the requirement for back propagation of errors: If the output synapses of a neuron undergo LTP/LTD, then its input synapses undergo similar changes. However, it is not clear whether the spreading LTP/LTD can be integrated by the neuron, analogous to the linear summation of error signals in the back-propagation algorithm. In addition, the spread

of LTP/LTD in culture does not appear to go beyond the immediate layer of neurons involved in the induction of LTP/LTD.

The Unit of Synaptic Modification The exclusively "presynaptic" pattern of LTP spread in culture (Tao et al 2000) suggests that many output synapses of a single excitatory neuron may be modified as an integral unit in activity-dependent synaptic plasticity, with changes at each synapse affecting the others. Such coordinated changes among specific sets of synapses may help to recruit new cells into an existing ensemble and contribute to the development of synchronous firing in the cortex. Furthermore, both the induction and the presynaptic spread of LTP are target cell specific (do not occur at glutamatergic synapses made onto GABAergic postsynaptic neurons), whereas GABAergic synapses may be modified directly by correlated activity and indirectly by spread of LTP/LTD from neighboring sites (Fitzsimonds et al 1997, Aizenman et al 1998, Tao et al 2000; K Ganguly, M-m Poo, unpublished results). Target cell–specific induction and spread of LTP/LTD indicate that the unit of synaptic modification is the subset of presynaptic nerve terminals innervating the same postsynaptic cell type. The ability of GABAergic synapses to be modified suggests an active role for these inhibitory connections in the learning and memory function of neural circuits. Finally, we note that for converging inputs in networks of dissociated hippocampal neurons in culture and in CA1 pyramidal neurons in hippocampal slices, input specificity appears to be more strictly followed for LTP than for LTD. The implication of such asymmetry in input specificity in the development and functioning of neural network remains to be investigated.

CONCLUDING REMARKS

The induction and expression of LTP/LTD represent only the initial consequence of correlated activity. A myriad of subsequent events, including local activation of enzymes, posttranslational modification of proteins, and changes in protein synthesis and gene expression, occur in both pre- and postsynaptic cells, leading to long-lasting structural and functional alterations of the synapses as well as the entire neuron (Milner et al 1998). The existence of long-range and selective signaling within the neuron, as examplified by the selective spread of LTP/LTD, argues that activity-induced local synaptic modification must be considered as an integral part of global neuronal changes. The existence of reciprocal and selective transsynaptic signaling between neurons further suggests that synaptic and neuronal plasticity must be considered within the context of the interacting neurons in the neural network.

Synaptic modification by correlated activity was proposed by Hebb as a mechanism for the formation of "cell assembly," a network consolidated by patterned activities, which serves as a fundamental unit for perceptual functions of the brain. Although Hebb's original notion that development of the network is intimately

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

linked to the perceptual function of the network remains to be validated, his postulate on synaptic modification has become a recurrent theme for understanding the synaptic basis underlying both activity-driven refinement of developing networks and learning and memory functions of mature networks. Developmental refinement may indeed use synaptic mechanisms similar to that used for learning and memory. In the spirit of Hebb, one may further ask whether the development of new networks is a basis of learning and memory in mature brain and, conversely, whether the structure of mature brain is but a reflection of perceptual experience during its early development.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Visit the Annual Reviews home page at www.AnnualReviews.org

LITERATURE CITED

Abbott LF, Blum KI. 1996. Functional significance of long-term potentiation for sequence learning and prediction, *Cereb. Cortex* 6:406–16

Abbott LF, Song S. 1999. Temporally asymmetric Hebbian learning, spike timing and neuronal response variability. In *Advances in Neural Information Processing Systems*, ed. MS Kearns, SA Solla, DA Cohn, 11:69–75. Cambridge, MA: MIT Press

Abeles M. 1991. *Corticonics*. Cambridge, UK: Cambridge Univ. Press

Abeles M, Bergman H, Margalit E, Vaadia E. 1993. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. Neurophysiol* 70:1629–38

Abraham WC, Goddard GV. 1983. Asymmetric relationships between homosynaptic long-term potentiation and heterosynaptic long-term depression. *Nature* 305:717–19

Abrams TW, Karl KA, Kandel ER. 1991. Biochemical studies of stimulus convergence during classical conditioning in Aplysia: dual regulation of adenylate cyclase by Ca^{2+}/calmodulin and transmitter. *J. Neurosci.* 11:2655–65

Aizenman CD, Manis PB, Linden DJ. 1998. Polarity of long-term synaptic gain change is related to postsynaptic spike firing at a cerebellar inhibitory synapse. *Neuron* 21:827–35

Andersen P, Sundberg SH, Sveen O, Wigström H. 1977. Specific long-lasting potentiation of synaptic transmission in hippocampal slices. *Nature* 266:736–37

Barrionuevo G, Brown TH. 1983. Associative long-term potentiation in hippocampal slices. *Proc. Natl. Acad. Sci. USA* 80:7347–51

Bear MF, Abraham WC. 1996. Long-term depression in hippocampus. *Annu. Rev. Neurosci.* 19:437–62

Bekkers JM, Richerson GB, Stevens CF. 1990. Origin of variability in quantal size in cultured hippocampal neurons and hippocampal slices. *Proc. Natl. Acad. Sci. USA* 87:5359–62

Bell CC, Han VZ, Sugawara Y, Grant K. 1997. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387:278–81

Berridge MJ. 1998. Neuronal calcium signaling. *Neuron* 21:13–26

Bi G-q, Morris RL, Liao GC, Alderton JM, Scholey JM, Steinhardt RA. 1997. Kinesin- and myosin-driven steps of vesicle recruitment for Ca^{2+}-regulated exocytosis. *J. Cell Biol.* 138:999–1008

Bi G-q, Poo M-m. 1998. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18:10464–72

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Bi G-q, Poo M-m. 1999. Distributed synaptic modification in neural networks induced by patterned stimulation. *Nature* 401:792–96

Bienenstock EL, Cooper LN, Munro PW. 1982. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* 2:32–48

Bliss TV, Collingridge GL. 1993. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361:31–39

Bliss TV, Gardner-Medwin AR. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the unanaestetized rabbit following stimulation of the perforant path. *J. Physiol.* 232:357–74

Bliss TV, Lømo T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232:331–56

Bonhoeffer T, Staiger V, Aertsen A. 1989. Synaptic plasticity in rat hippocampal slice cultures: local "Hebbian" conjunction of pre- and postsynaptic stimulation leads to distributed synaptic enhancement. *Proc. Natl. Acad. Sci. USA* 86:8113–17

Bortolotto ZA, Fitzjohn SM, Collingridge GL. 1999. Roles of metabotropic glutamate receptors in LTP and LTD in the hippocampus. *Curr. Opin. Neurobiol.* 9:299–304

Bourne HR, Nicoll R. 1993. Molecular machines integrate coincident synaptic signals. *Cell* 72(Suppl.):65–75

Bradler JE, Barrioneuvo G. 1989. Long-term potentiation in hippocampal CA3 neurons: tetanized input regulates heterosynaptic efficacy. *Synapse* 4:132–42

Braitenberg V. 1967. Is the cerebellar cortex a biological clock in the millisecond range? *Prog. Brain Res.* 25:334–46

Brown TH, Kairiss EW, Keenan CL. 1990. Hebbian synapses: biophysical mechanisms and algorithms *Annu. Rev. Neurosci.* 13:475–511

Bullock TH. 1970. The reliability of neurons. *J. Gen. Physiol.* 55:565–84

Buonomano DV, Hickmott PW, Merzenich MM. 1997. Context-sensitive synaptic plasticity and temporal-to-spatial transformations in hippocampal slices. *Proc. Natl. Acad. Sci. USA* 94:10403–8

Buonomano DV, Merzenich MM. 1995. Temporal information transformed into a spatial code by a neural network with realistic properties. *Science* 267:1028–30

Buonomano DV, Merzenich MM. 1998. Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* 21:149–86

Buzsáki G, Penttonen M, Nádasdy Z, Bragin A. 1996. Pattern and inhibition-dependent invasion of pyramidal cell dendrites by fast spikes in the hippocampus in vivo. *Proc. Natl. Acad. Sci. USA* 93:9921–25

Calvin WH, Stevens CF. 1968. Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J. Neurophysiol.* 31:574–87

Carew TJ, Hawkins RD, Abrams TW, Kandel ER. 1984. A test of Hebb's postulate at identified synapses which mediate classical conditioning in Aplysia. *J. Neurosci.* 4:1217–24

Carew TJ, Walters ET, Kandel ER. 1981. Classical conditioning in a simple withdrawal reflex in *Aplysia californica. J. Neurosci.* 1:1426–37

Carr CE. 1993. Processing of temporal information in the brain. *Annu. Rev. Neurosci.* 16:223–43

Cash S, Zucker RS, Poo M-m. 1996. Spread of synaptic depression mediated by presynaptic cytoplasmic signaling. *Science* 272:998–1001

Clark GA, Hawkins RD, Kandel ER. 1994. Activity-dependent enhancement of presynaptic facilitation provides a cellular mechanism for the temporal specificity of classical conditioning in Aplysia. *Learn. Mem.* 1:243–57

Clark GA, Kandel ER. 1993. Induction of long-term facilitation in Aplysia sensory neurons by local application of serotonin to remote synapses. *Proc. Natl. Acad. Sci. USA* 90:11411–15

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Collingridge GL, Kehl SJ, McLennan H. 1983. Excitatory amino acids in synaptic transmission in the Schaffer collateral-commissural pathway of the rat hippocampus. *J. Physiol.* 334:33–46

Constantine-Paton M, Cline HT, Debski E. 1990. Patterned activity, synaptic convergence, and the NMDA receptor in developing visual pathways. *Annu. Rev. Neurosci.* 13:129–54

Cowan AI, Stricker C, Reece LJ, Redman SJ. 1998. Long-term plasticity at excitatory synapses on aspinous interneurons in area CA1 lacks synaptic specificity. *J. Neurophysiol.* 79:13–20

Crair MC, Gillespie DC, Stryker MP. 1998. The role of visual experience in the development of columns in cat visual cortex. *Science* 279:566–70

Crowley JC, Katz LC. 1999. Development of ocular dominance columns in the absence of retinal input. *Nat. Neurosci.* 2:1125–30

Dan Y, Alonso JM, Usrey WM, Reid RC. 1998. Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nat. Neurosci.* 1:501–7

Debanne D, Gahwiler BH, Thompson SM. 1998. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol.* 507:237–47

deCharms RC, Merzenich MM. 1996. Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* 381:610–13

Diesmann M, Gewaltig MO, Aertsen A. 1999. Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–33

Dudek SM, Bear MF. 1992. Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proc. Natl. Acad. Sci. USA* 89:4363–67

Dunwiddie T, Lynch G. 1978. Long-term potentiation and depression of synaptic responses in the rat hippocampus: localization and frequency dependency. *J. Physiol.* 276:353–67

Egger V, Feldmeyer D, Sakmann B. 1999. Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat. Neurosci.* 2:1098–105

Emptage NJ, Carew TJ. 1993. Long-term synaptic facilitation in the absence of short-term facilitation in Aplysia neurons. *Science* 262:253–56

Engert F, Bonhoeffer T. 1997. Synapse specificity of long-term potentiation breaks down at short distances. *Nature* 388:279–84

Erwin E, Miller KD. 1998. Correlation-based development of ocularly matched orientation and ocular dominance maps: determination of required input activities. *J. Neurosci.* 18:9870–95

Fitzsimonds RM, Poo M-m. 1998. Retrograde signaling in the development and modification of synapses. *Physiol. Rev.* 78:143–70

Fitzsimonds RM, Song H-j, Poo M-m. 1997. Propagation of activity-dependent synaptic depression in simple neural networks. *Nature* 388:439–48

Fregnac Y, Bienenstock E. 1998. Correlational models of synaptic plasticity: development, learning, and cortical dynamics of mental representations. In *Mechanistic Relationships Between Development and Learning*, ed. T Carew, R Menzel, C Shatz. pp. 113–48. New York: Wiley

Gerstner W, Abbott LF. 1997. Learning navigational maps through potentiation and modulation of hippocampal place cells. *J. Comput. Neurosci.* 4:79–94

Gerstner W, Kempter R, van Hemmen JL, Wagner H. 1996. A neuronal learning rule for sub-millisecond temporal coding. *Nature* 383:76–81

Goda Y, Stevens CF. 1996. Long-term depression properties in a simple system. *Neuron* 16:103–11

Gustafsson B, Wigström H. 1986. Hippocampal long-lasting potentiation produced by pairing single volleys and brief conditioning tetani evoked in separate afferents. *J. Neurosci.* 6:1575–82

Gustafsson B, Wigström H, Abraham WC, Huang YY. 1987. Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.* 7:774–80

Hawkins RD, Abrams TW, Carew TJ, Kandel ER. 1983. A cellular mechanism of classical conditioning in Aplysia: activity-dependent amplification of presynaptic facilitation. *Science* 219:400–5

Hebb D. 1949. The *Organization of Behavior*. New York: Wiley

Hempel CM, Vincent P, Adams SR, Tsien RY, Selverston AI. 1996. Spatio-temporal dynamics of cyclic AMP signals in an intact neural circuit. *Nature* 384:166–69

Higashima M, Yamamoto C. 1985. Two components of long-term potentiation in mossy fiber-induced excitation in hippocampus. *Exp. Neurol.* 90:529–39

Hirsch JC, Barrionuevo G, Crepel F. 1992. Homo- and heterosynaptic changes in efficacy are expressed in prefrontal neurons: an in vitro study in the rat. *Synapse* 12:82–85

Hoffman DA, Magee JC, Colbert CM, Johnston D. 1997. K+ channel regulation of signal propagation in dendrites of hippocampal pyramidal neurons. *Nature* 387:869–75

Hollmann M, Heinemann S. 1994. Cloned glutamate receptors. *Annu. Rev. Neurosci.* 17:31–108

Hopfield JJ. 1995. Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–36

Horton JC, Hocking DR. 1996. An adult-like pattern of ocular dominance columns in striate cortex of newborn monkeys prior to visual experience. *J. Neurosci.* 16:1791–807

Hubel DH, Wiesel TN. 1965. Binocular interaction in striate cortex of kittens reared with artificial squint. *J. Neurophysiol.* 28:1041–59

Ito M, Sakurai M, Tongroach P. 1982. Climbing fibre induced depression of both mossy fibre responsiveness and glutamate sensitivity of cerebellar Purkinje cells. *J. Physiol.* 324:113–34

Jahr CE, Stevens CF. 1987. Glutamate activates multiple single channel conductances in hippocampal neurons. *Nature* 325:522–5

Jeffress LA. 1948. A place theory of sound localization. *J. Comp. Physiol. Psychol.* 41:35–39

Johnston D, Magee JC, Colbert CM, Cristie BR. 1996. Active properties of neuronal dendrites. *Annu. Rev. Neurosci.* 19:165–86

Katz LC, Shatz CJ. 1996. Synaptic activity and the construction of cortical circuits. *Science* 274:1133–38

Kelso SR, Brown TH. 1986. Differential conditioning of associative synaptic enhancement in hippocampal brain slices. *Science* 232:85–87

Kelso SR, Ganong AH, Brown TH. 1986. Hebbian synapses in hippocampus. *Proc. Natl. Acad. Sci. USA* 83:5326–30

Kempter R, Gerstner W, van Hemmen JL. 2000. Hebbian learning and spiking neurons. *Phys. Rev. E* 59. In press

Koester HJ, Sakmann B. 1998. Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials *Proc. Natl. Acad. Sci. USA* 95:9596–601

Kossel A, Bonhoeffer T, Bolz J. 1990. Non-Hebbian synapses in rat visual cortex. *NeuroReport* 1:115–18

Kuznetsov SA, Langford GM, Weiss DG. 1992. Actin-dependent organelle movement in squid axoplasm. *Nature* 356:722–25

Lasser-Ross N, Ross WN. 1992. Imaging voltage and synaptically activated sodium transients in cerebellar Purkinje cells. *Proc. R. Soc. London Ser. B* 247:35–39

Laurent G. 1997. Olfactory processing: maps, time and codes. *Curr. Opin. Neurobiol.* 7:547–53

Lechner HA, Byrne JH. 1998. New perspectives on classical conditioning: a synthesis of Hebbian and non-Hebbian mechanisms. *Neuron* 20:355–58

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Levy WB, Steward O. 1979. Synapses as associative memory elements in the hippocampal formation. *Brain Res.* 175:233–45

Levy WB, Steward O. 1983. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8:791–97

Lin XY, Glanzman DL. 1994. Hebbian induction of long-term potentiation of Aplysia sensorimotor synapses: partial requirement for activation of an NMDA-related receptor. *Proc. R. Soc. London Ser. B* 255:215–21

Linden DJ. 1999. The return of the spike: postsynaptic action potentials and the induction of LTP and LTD. *Neuron* 22:661–66

Linden DJ, Connor JA. 1995. Long-term synaptic depression. *Annu. Rev. Neurosci.* 18:319–57

Lisman J. 1989. A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA* 86:9574–78

Lo YJ, Poo MM. 1991. Activity-dependent synaptic competition in vitro: heterosynaptic suppression of developing synapses. *Science* 254:1019–22

Lux HD, Veselovsky NS. 1994. Glutamate-produced long-term potentiation by selective challenge of presynaptic neurons in rat hippocampal cultures. *Neurosci. Lett.* 178:231–34

Lynch GS, Dunwiddie T, Gribkoff V. 1977. Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature* 266:737–39

Maass W, Bishop CM, eds. 1999. *Pulsed Neural Network.* Cambridge, MA: MIT Press. 377 pp.

Maccaferri G, Toth K, McBain CJ. 1998. Target-specific expression of presynaptic mossy fiber plasticity. *Science* 279:1368–70

Magee JC, Johnston D. 1997. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275:209–13

Mainen ZF, Sejnowski TJ. 1995. Reliability of spike timing in neocortical neurons. *Science* 268:1503–6

Malenka RC, Kauer JA, Perkel DJ, Mauk MD, Kelly PT, et al. 1989. An essential role for postsynaptic calmodulin and protein kinase activity in long-term potentiation. *Nature* 340:554–57

Malenka RC, Nicoll RA. 1999. Long-term potentiation—a decade of progress. *Science* 285:1870–74

Malinow R, Schulman H, Tsien RW. 1989. Inhibition of postsynaptic PKC or CaMKII blocks induction but not expression of LTP. *Science* 245:862–66

Markram H, Lubke J, Frotscher M, Sakmann B. 1997. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275:213–15

Marsálek P, Koch C, Maunsell J. 1997. On the relationship between synaptic input and spike output jitter in individual neurons. *Proc. Natl. Acad. Sci. USA* 94:735–40

Mayer ML, Westbrook GL, Guthrie PB. 1984. Voltage-dependent block by Mg^{2+} of NMDA responses in spinal cord neurones. *Nature* 309:261–63

McMahon LL, Kauer JA. 1997. Hippocampal interneurons express a novel form of synaptic plasticity. *Neuron* 18:295–305

McNaughton BL, Douglas RM, Goddard GV. 1978. Synaptic enhancement in fascia dentata: cooperativity among coactive afferents. *Brain Res.* 157:277–93

Mehta MR, Barnes CA, McNaughton BL. 1997. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* 94:8918–21

Mehta MR, Quirk MC, Wilson MA. 2000. Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25:707–15

Miller KD. 1994. A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *J. Neurosci.* 14:409–41

Miller KD, Keller JB, Stryker MP. 1989. Ocular dominance column development: analysis and simulation. *Science* 245:605–15

Milner B, Squire LR, Kandel ER. 1998. Cognitive neuroscience and the study of memory. *Neuron* 20:445–68

Minai AA, Levy WB. 1993. Sequence learning in a single trial. *INNS World Congr. Neural Netw.* 2:505–8

Misgeld U, Sarvey JM, Klee MR. 1979. Heterosynaptic postactivation potentiation in hippocampal CA 3 neurons: long-term changes of the postsynaptic potentials. *Exp. Brain Res.* 37:217–29

Montague PR, Sejnowski TJ. 1994. The predictive brain: temporal coincidence and temporal order in synaptic learning mechanisms *Learn. Mem.* 1:1–33

Moore JW, Choi J-S, Brunzell DH. 1998. Predictive timing under temporal uncertainty: the time derivative model of the conditioned response. In *Timing of Behavior*, ed. DA Risenbaum, CE Collyer, pp. 3–34. Cambridge, MA: MIT Press

Mulkey RM, Endo S, Shenolikar S, Malenka RC. 1994. Involvement of a calcineurin/inhibitor-1 phosphatase cascade in hippocampal long-term depression. *Nature* 369:486–88

Mulkey RM, Malenka RC. 1992. Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus. *Neuron* 9:967–75

Muller D, Hefft S, Figurov A. 1995. Heterosynaptic interactions between LTP and LTD in CA1 hippocampal slices. *Neuron* 14:599–605

Murphy GG, Glanzman DL. 1997. Mediation of classical conditioning in *Aplysia californica* by long-term potentiation of sensorimotor synapses. *Science* 278:467–71

Nakamura T, Barbara JG, Nakamura K, Ross WN. 1999. Synergistic release of Ca^{2+} from IP3-sensitive stores evoked by synaptic activation of mGluRs paired with backpropagating action potentials. *Neuron* 24:727–37

Nicoll RA, Malenka RC. 1997. Neurobiology. Long-distance long-term depression. *Nature* 388:427–28

Nishiyama M, Hong K, Mikoshiba K, Poo M-m, Kato K. 2000. Calcium release from internal stores regulates polarity and input specificity of synaptic modification. *Nature.* In press

Nowak L, Bregestovski P, Ascher P, Herbet A, Prochiantz A. 1984. Magnesium gates glutamate-activated channels in mouse central neurones. *Nature* 307:462–65

Paulsen O, Sejnowski TJ. 2000. Natural patterns of activity and long-term synaptic plasticity. *Curr. Opin. Neurobiol.* 10:172–79

Penn AA, Riquelme PA, Feller MB, Shatz CJ. 1998. Competition in retinogeniculate patterning driven by spontaneous activity. *Science* 279:2108–12

Perkel DH, Bullock TH. 1968. Neural coding. *Neurosci. Res. Prog. Sum.* 3:405–527

Purves D. 1975. Functional and structural changes in mammalian sympathetic neurones following interruption of their axons. *J. Physiol.* 252:429–63

Rao RPN, Sejnowski TJ. 2000. Predictive learning of temporal sequences in recurrent neocortical circuits In *Advances in Neural Information Processing Systems*, Vol. 12, ed. SA Solla, TK Leen, K-R Muller. Cambridge, MA: MIT Press. In press

Regehr WG, Konnerth A, Armstrong CM. 1992. Sodium action potentials in the dendrites of cerebellar Purkinje cells. *Proc. Natl. Acad. Sci. USA* 89:5492–96

Reyes A, Lujan R, Burnashev N, Somogyi P, Sakmann B. 1998. Target-cell-specific facilitation and depression in neocortical circuits. *Nat. Neurosci.* 1:279–85

Riehle A, Grün S, Diesmann M, Aertsen A. 1997. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950–53

Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W. 1997. *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press. 395 pp.

Roberts PD. 1999. Computational consequences of temporally asymmetric learning

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

rules. I. Differential hebbian learning. *J. Comput. Neurosci.* 7:235–46

Roelfsema PR, Engel AK, König P, Singer W. 1997. Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature* 385:157–61

Rolls ET, Treves A. 1998. *Neural Networks and Brain Function.* New York: Oxford Univ. Press

Rumelhart DE, Hinton GE, Williams RJ. 1986a. Learning internal representations by error propagation. In *Parallel Distributed Processing*, ed. JA Feldman, PJ Hayes, DE Rumelhart. 1:318–62. Cambridge, MA: MIT Press

Rumelhart DE, Hinton GE, Williams RJ. 1986b. Learning representation by back-propagating errors. *Nature* 323:533–36

Sastry BR, Goh JW, Auyeung A. 1986. Associative induction of posttetanic and long-term potentiation in CA1 neurons of rat hippocampus. *Science* 232:988–90

Schinder AF, Berninger B, Poo M. 2000. Postsynaptic target specificity of neurotrophin-induced presynaptic potentiation. *Neuron* 25:151–63

Schuman EM, Madison DV. 1994. Locally distributed synaptic potentiation in the hippocampus. *Science* 263:532–36

Sejnowski TJ. 1977a. Statistical constraints on synaptic plasticity. *J. Theor. Biol.* 69:385–89

Sejnowski TJ. 1977b. Storing covariance with nonlinearly interacting neurons. *J. Math. Biol.* 4:303–21

Sejnowski TJ. 1999. The book of Hebb. *Neuron* 24:773–76

Sejnowski TJ, Tesauro G. 1989. The Hebb rule for synaptic plasticity: algorithms and implementations. In *Neural Models of Plasticity: Experimental and Theoretical Approaches*, ed. John H. Byrne. pp. 94–103. San Diego, CA: Academic

Senn W, Markram H, Tsodyks M. 2000. An algorithm for modifying neurotransmitter release probability based on pre- and postsynaptic spike timing. *Neural Comput.* In press

Silva AJ, Stevens CF, Tonegawa S, Wang Y. 1992. Deficient hippocampal long-term potentiation in alpha-calcium-calmodulin kinase II mutant mice. *Science* 257:201–6

Singer W. 1999. Time as coding space? *Curr. Opin. Neurobiol.* 9:189–94

Singer W, Gray CM. 1995. Visual feature integration and the temporal correlation hypothesis *Annu. Rev. Neurosci.* 18:555–86

Squire LR. 1987. *Memory and Brain.* New York: Oxford Univ. Press

Staubli UV, Ji ZX. 1996. The induction of homo- vs. heterosynaptic LTD in area CA1 of hippocampal slices from adult rats. *Brain Res.* 714:169–76

Stent GS. 1973. A physiological mechanism for Hebb's postulate of learning. *Proc. Natl. Acad. Sci. USA* 70:997–1001

Stevens CF. 1996. Strengths and weaknesses in memory. *Nature* 381:471–72

Stevens CF, Zador AM. 1998. Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.* 1:210–17

Stryker MP, Strickland SL. 1984. Physiological segregation of ocular dominance columns depends on patterns of afferent activity. *Invest. Ophthalmol. Suppl.* 25:278

Stuart GJ, Sakmann B. 1994. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367:69–72

Sutton RS, Barto AG. 1981. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol. Rev.* 88:135–70

Tank DW, Hopfield JJ. 1987. Neural computation by concentrating information in time. *Proc. Natl. Acad. Sci. USA* 84:1896–900

Tao HW, Zhang LI, Bi G-q, Poo M-m. 2000. Selective presynaptic propagation of long-term potentiation in defined neural networks. *J. Neurosci.* 20:3233–43

Tesauro G. 1986. Simple neural models of classical conditioning. *Biol. Cybern.* 55:187–200

Teyler TJ, DiScenna P. 1987. Long-term potentiation. *Annu. Rev. Neurosci.* 10:131–61

Thompson RF, Krupa DJ. 1994. Organization of memory traces in the mammalian brain. *Annu. Rev. Neurosci.* 17:519–49

Tsumoto T, Suda K. 1979. Cross-depression: an electrophysiological manifestation of binocular competition in the developing visual cortex. *Brain Res.* 168:190–94

Tzounopoulos T, Janz R, Südhof TC, Nicoll RA, Malenka RC. 1998. A role for cAMP in long-term depression at hippocampal mossy fiber synapses. *Neuron* 21:837–45

Vale RD, Reese TS, Sheetz MP. 1985. Identification of a novel force-generating protein, kinesin, involved in microtubule-based motility. *Cell* 42:39–50

Vallee RB, Bloom GS. 1991. Mechanisms of fast and slow axonal transport. *Annu. Rev. Neurosci.* 14:59–92

Walters ET, Byrne JH. 1983. Associative conditioning of single sensory neurons suggests a cellular mechanism for learning. *Science* 219:405–8

Weliky M, Katz LC. 1997. Disruption of orientation tuning in visual cortex by artificially correlated neuronal activity. *Nature* 386:680–85

Wiesel TN, Hubel DH. 1963. Single cell responses in striate cortex of kittens deprived of vision in one eye. *J. Neurophysiol.* 26:1003–7

Wigström H, Gustafsson B, Huang YY, Abraham WC. 1986. Hippocampal long- term potentiation is induced by pairing single afferent volleys with intracellularly injected depolarizing current pulses. *Acta Physiol. Scand.* 126:317–19

Wong RO, Chernjavsky A, Smith SJ, Shatz CJ. 1995. Early functional neural networks in the developing retina. *Nature* 374:716–18

Yamamoto C, Chujo T. 1978. Long-term potentiation in thin hippocampal sections studied by intracellular and extracellular recordings. *Exp. Neurol.* 58:242–50

Yang SN, Tang YG, Zucker RS. 1999. Selective induction of LTP and LTD by postsynaptic $[Ca^{2+}]$ elevation. *J. Neurophysiol.* 81:781–87

Yeckel MF, Kapur A, Johnston D. 1999. Multiple forms of LTP in hippocampal CA3 neurons use a common postsynaptic mechanism. *Nat. Neurosci.* 2:625–33

Yovell Y, Abrams TW. 1992. Temporal asymmetry in activation of Aplysia adenylyl cyclase by calcium and transmitter may explain temporal requirements of conditioning. *Proc. Natl. Acad. Sci. USA* 89:6526–30

Yovell Y, Kandel ER, Dudai Y, Abrams TW. 1992. A quantitative study of the Ca^{2+}/calmodulin sensitivity of adenylyl cyclase in Aplysia, Drosophila, and rat. *J. Neurochem.* 59:1736–44

Zhang LI, Tao HW, Holt CE, Harris WA, Poo M-m. 1998. A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395:37–44

Zheng JQ, Zheng Z, Poo M. 1994. Long-range signaling in growing neurons after local elevation of cyclic AMP-dependent activity. *J. Cell Biol.* 127:1693–701

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annual Review of Neuroscience
Volume 24, 2001

CONTENTS

PDZ DOMAINS AND THE ORGANIZATION OF SUPRAMOLECULAR COMPLEXES, *Morgan Sheng, Carlo Sala* 1

THE ROLE AND REGULATION OF ADENOSINE IN THE CENTRAL NERVOUS SYSTEM, *Thomas V. Dunwiddie, Susan A. Masino* 31

LOCALIZATION AND GLOBALIZATION IN CONSCIOUS VISION, *S. Zeki* 57

GLIAL CONTROL OF NEURONAL DEVELOPMENT, *Greg Lemke* 87

TOUCH AND GO: Decision-Making Mechanisms in Somatosensation, *Ranulfo Romo, Emilio Salinas* 107

SYNAPTIC MODIFICATION BY CORRELATED ACTIVITY: Hebb"s Postulate Revisited, *Guo-qiang Bi, Mu-ming Poo* 139

AN INTEGRATIVE THEORY OF PREFRONTAL CORTEX FUNCTION, *Earl K. Miller, Jonathan D. Cohen* 167

THE PHYSIOLOGY OF STEREOPSIS, *B. G. Cumming, G. C. DeAngelis* 203

PARANEOPLASTIC NEUROLOGIC DISEASE ANTIGENS: RNA-Binding Proteins and Signaling Proteins in Neuronal Degeneration, *Kiran Musunuru, Robert B. Darnell* 239

ODOR ENCODING AS AN ACTIVE, DYNAMICAL PROCESS: Experiments, Computation, and Theory, *Gilles Laurent, Mark Stopfer, Rainer W Friedrich, Misha I Rabinovich, Alexander Volkovskii, Henry DI Abarbanel* 263

PROTEIN SYNTHESIS AT SYNAPTIC SITES ON DENDRITES, *Oswald Steward, Erin M. Schuman* 299

SIGNALING AND TRANSCRIPTIONAL MECHANISMS IN PITUITARY DEVELOPMENT, *Jeremy S. Dasen, Michael G. Rosenfeld* 327

NEUROPEPTIDES AND THE INTEGRATION OF MOTOR RESPONSES TO DEHYDRATION , *Alan G. Watts* 357

THE DEVELOPMENTAL BIOLOGY OF BRAIN TUMORS, *Robert Wechsler-Reya, Matthew P. Scott* 385

TO EAT OR TO SLEEP? OREXIN IN THE REGULATION OF FEEDING AND WAKEFULNESS, *Jon T. Willie, Richard M. Chemelli, Christopher M. Sinton, Masashi Yanagisawa* 429

SPATIAL PROCESSING IN THE BRAIN: The Activity of Hippocampal Place Cells, *Phillip J. Best, Aaron M. White, Ali Minai* 459

THE VANILLOID RECEPTOR: A Molecular Gateway to the Pain Pathway, *Michael J Caterina, David Julius* 487

PRION DISEASES OF HUMANS AND ANIMALS: Their Causes and Molecular Basis, *John Collinge* 519

VIKTOR HAMBURGER AND RITA LEVI-MONTALCINI: The Path to the Discovery of Nerve Growth Factor, *W. Maxwell Cowan* 551

EARLY DAYS OF THE NERVE GROWTH FACTOR PROTEINS, *Eric M. Shooter* 601

SEQUENTIAL ORGANIZATION OF MULTIPLE MOVEMENTS: Involvement of Cortical Motor Areas, *Jun Tanji* 631

INFLUENCE OF DENDRITIC CONDUCTANCES ON THE INPUT-OUTPUT PROPERTIES OF NEURONS, *Alex Reyes* 653

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

NEUROTROPHINS: Roles in Neuronal Development and Function, *Eric J Huang, Louis F Reichardt* 677
CONTRIBUTIONS OF THE MEDULLARY RAPHE AND VENTROMEDIAL RETICULAR REGION TO PAIN MODULATION AND OTHER HOMEOSTATIC FUNCTIONS, *Peggy Mason* 737
ACTIVATION, DEACTIVATION, AND ADAPTATION IN VERTEBRATE PHOTORECEPTOR CELLS, *Marie E Burns, Denis A Baylor* 779
ACTIVITY-DEPENDENT SPINAL CORD PLASTICITY IN HEALTH AND DISEASE, *Jonathan R Wolpaw, Ann M Tennissen* 807
QUANTITATIVE GENETICS AND MOUSE BEHAVIOR, *Jeanne M Wehner, Richard A Radcliffe, Barbara J Bowers* 845
EARLY ANTERIOR/POSTERIOR PATTERNING OF THE MIDBRAIN AND CEREBELLUM, *Aimin Liu, Alexandra L Joyner* 869
NEUROBIOLOGY OF PAVLOVIAN FEAR CONDITIONING, *Stephen Maren* 897
{{alpha}}-LATROTOXIN AND ITS RECEPTORS: Neurexins and CIRL/Latrophilins, *Thomas C Südhof* 933
IMAGING AND CODING IN THE OLFACTORY SYSTEM, *John S Kauer, Joel White* 963
THE ROLE OF THE CEREBELLUM IN VOLUNTARY EYE MOVEMENTS, *Farrel R Robinson, Albert F Fuchs* 981
ROLE OF THE REELIN SIGNALING PATHWAY IN CENTRAL NERVOUS SYSTEM DEVELOPMENT, *Dennis S Rice, Tom Curran* 1005
HUMAN BRAIN MALFORMATIONS AND THEIR LESSONS FOR NEURONAL MIGRATION, *M Elizabeth Ross, Christopher A Walsh* 1041
MORPHOLOGICAL CHANGES IN DENDRITIC SPINES ASSOCIATED WITH LONG-TERM SYNAPTIC PLASTICITY, *Rafael Yuste, Tobias Bonhoeffer* 1071

STOPPING TIME: The Genetics of Fly and Mouse Circadian Clocks, *Ravi Allada, Patrick Emery, Joseph S. Takahashi, Michael Rosbash* 1091
NEURODEGENERATIVE TAUOPATHIES, *Virginia M-Y Lee, Michel Goedert, John Q Trojanowski* 1121
MATERNAL CARE, GENE EXPRESSION, AND THE TRANSMISSION OF INDIVIDUAL DIFFERENCES IN STRESS REACTIVITY ACROSS GENERATIONS, *Michael J Meaney* 1161
NATURAL IMAGE STATISTICS AND NEURAL REPRESENTATION, *Eero P Simoncelli, Bruno A Olshausen* 1193
Nerve Growth Factor Signaling, Neuroprotection, and Neural Repair, *Michael V Sofroniew, Charles L Howe, William C Mobley* 1217
FLIES, GENES, AND LEARNING, *Scott Waddell, William G Quinn* 1283

Annu. Rev. Neurosci. 2001.24:139-166. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

ANNUAL REVIEWS **Further**

Click here for quick links to Annual Reviews content online, including:

- Other articles in this volume
- Top cited articles
- Top downloaded articles
- Our comprehensive search

Spike Timing–Dependent Plasticity: A Hebbian Learning Rule

Natalia Caporale and Yang Dan

Division of Neurobiology, Department of Molecular and Cell Biology, and Helen Wills Neuroscience Institute, University of California, Berkeley, California 94720; email: caporale@socrates.berkeley.edu, ydan@berkeley.edu

Annu. Rev. Neurosci. 2008. 31:25–46

First published online as a Review in Advance on February 14, 2008

The *Annual Review of Neuroscience* is online at neuro.annualreviews.org

This article's doi: 10.1146/annurev.neuro.31.060407.125639

Copyright © 2008 by Annual Reviews. All rights reserved

0147-006X/08/0721-0025$20.00

Key Words

long-term potentiation, long-term depression, synapse, memory, backpropagating action potential

Abstract

Spike timing–dependent plasticity (STDP) as a Hebbian synaptic learning rule has been demonstrated in various neural circuits over a wide spectrum of species, from insects to humans. The dependence of synaptic modification on the order of pre- and postsynaptic spiking within a critical window of tens of milliseconds has profound functional implications. Over the past decade, significant progress has been made in understanding the cellular mechanisms of STDP at both excitatory and inhibitory synapses and of the associated changes in neuronal excitability and synaptic integration. Beyond the basic asymmetric window, recent studies have also revealed several layers of complexity in STDP, including its dependence on dendritic location, the nonlinear integration of synaptic modification induced by complex spike trains, and the modulation of STDP by inhibitory and neuromodulatory inputs. Finally, the functional consequences of STDP have been examined directly in an increasing number of neural circuits in vivo.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Contents

INTRODUCTION 26
CELLULAR MECHANISMS 27
tLTP Window 28
tLTD Window 28
STDP OF INHIBITION 29
STDP of Excitatory Synapses onto Inhibitory Neurons 30
STDP of GABAergic Synapses 30
STDP WITH COMPLEX SPIKE PATTERNS 31
DEPENDENCE ON DENDRITIC LOCATION 32
MODULATION OF STDP BY OTHER INPUTS 34
PLASTICITY OF NEURONAL EXCITABILITY AND SYNAPTIC INTEGRATION 35
STDP IN VIVO 36
Electrical Stimulation 36
Paired Sensory Stimulation 37
Motion Stimuli 37
Sensory Deprivation 38
FINAL REMARKS 38

INTRODUCTION

Electrical activity plays crucial roles in the structural and functional refinement of neural circuits throughout an organism's lifetime (Buonomano & Merzenich 1998, Gilbert 1998, Karmarkar & Dan 2006, Katz & Shatz 1996). Manipulations of sensory experience that disrupt normal activity patterns can lead to large-scale network remodeling and marked changes in neural response properties. Learning and memory are also likely to be mediated by activity-dependent circuit modifications. Understanding the cellular mechanisms underlying such functional plasticity has been a long-standing challenge in neuroscience (Martin et al. 2000).

In his influential postulate on the cellular basis for learning, Hebb stated that "when an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased" (Hebb 1949). This postulate gained strong experimental support with the finding of long-term potentiation (LTP) of synaptic transmission, initially discovered in the hippocampus (Bliss & Gardner-Medwin 1973, Bliss & Lomo 1973) and subsequently reported in a large number of neural circuits, including various neocortical areas (Artola & Singer 1987, Iriki et al. 1989, Hirsch et al. 1992), the amygdala (Chapman et al. 1990, Clugnet & LeDoux 1990), and the midbrain reward circuit (Liu et al. 2005, Pu et al. 2006). Traditionally, LTP is induced by high-frequency stimulation (HFS) of the presynaptic afferents or by pairing low-frequency stimulation (LFS) with large postsynaptic depolarization (>30 mV). In contrast, long-term depression (LTD) is induced by LFS, either alone or paired with a small postsynaptic depolarization (Artola et al. 1990, Dudek & Bear 1993, Kirkwood & Bear 1994, Linden & Connor 1995, Mulkey & Malenka 1992, Stanton & Sejnowski 1989). Together, LTP and LTD allow activity-dependent bidirectional modification of synaptic strength, thus serving as promising candidates for the synaptic basis of learning and memory (Bliss & Collingridge 1993; Ito 2005; Siegelbaum & Kandel 1991).

To characterize the temporal requirements for the induction of LTP and LTD, Levi & Steward (1983) varied the relative timing of a strong and a weak input from the entorhinal cortex to the dental gyrus and found that synaptic modification depended on the temporal order of the two inputs. Potentiation was produced when the weak input preceded the strong input by less than 20 ms, and reversing the order led to depression. Subsequent studies further demonstrated the importance of the temporal order of pre- and postsynaptic spiking in synaptic modification and delineated the critical window on the order of tens of milliseconds (Bi & Poo 1998, Debanne et al. 1998, Magee & Johnston 1997, Markram et al. 1997,

LTP: long-term potentiation

HFS: high-frequency stimulation

LFS: low-frequency stimulation

LTD: long-term depression

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Zhang et al. 1998) (**Figure 1*a***, I). Such spike-timing-dependent plasticity (STDP) (Abbott & Nelson 2000) has now been observed at excitatory synapses in a wide variety of neural circuits (Boettiger & Doupe 2001, Cassenaer & Laurent 2007, Egger et al. 1999, Feldman 2000, Froemke & Dan 2002, Sjostrom et al. 2001, Tzounopoulos et al. 2004). Compared with the correlational forms of synaptic plasticity, STDP captures the importance of causality in determining the direction of synaptic modification, which is implied in Hebb's original postulate.

Recent studies have further characterized the mechanism and function of STDP in both in vitro and in vivo preparations, addressing the following questions: Which cellular mechanisms determine the STDP window, and how similar are they to the mechanisms underlying LTP and LTD induced by HFS and LFS, respectively? Does the window depend on the dendritic location of the input, and can it be regulated by neuromodulatory inputs? Does a similar learning rule apply to the inhibitory circuits? Can we observe the consequences of the asymmetric window in vivo, and can it account for the synaptic modifications induced by complex, naturalistic spike trains? In this review we summarize recent progress in these areas.

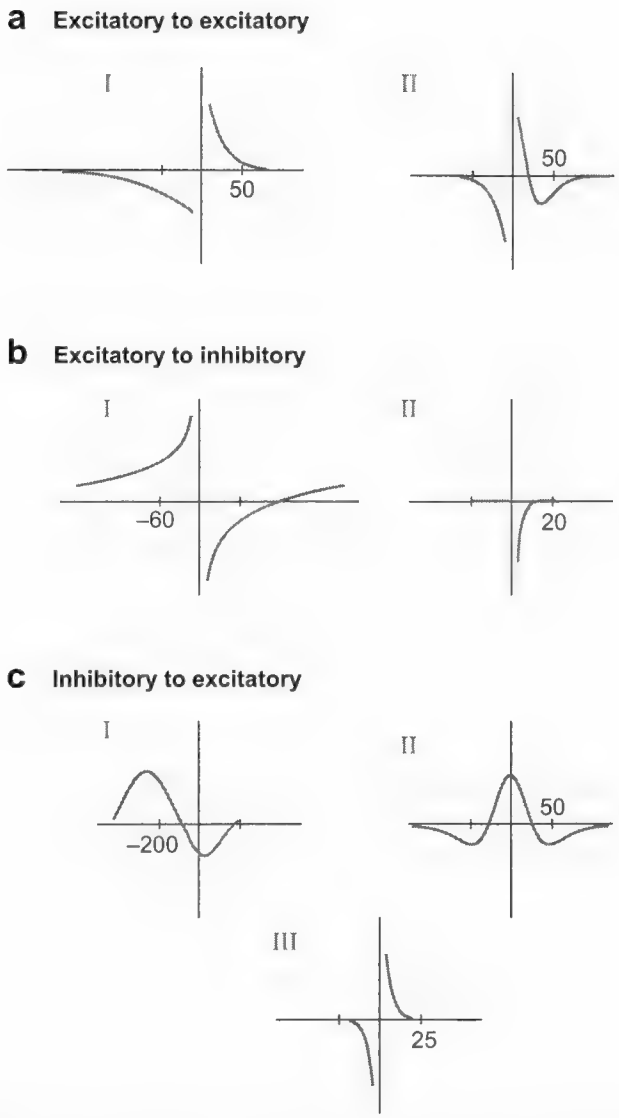

Figure 1

Diversity of temporal windows for STDP induction. *a:* Windows for excitatory to excitatory connections. *b:* Windows for excitatory to inhibitory connections. *c:* Windows for inhibitory to excitatory connections. Temporal axis is in milliseconds.

CELLULAR MECHANISMS

For many glutamatergic synapses, the inductions of LTP by HFS and LTD by LFS both require the activation of NMDA (N-methyl-d-aspartate) receptors and a rise in postsynaptic Ca^{2+} level (Malenka & Bear 2004). The NMDA receptor is thought to serve as the coincidence detector: The presynaptic activation provides glutamate and the postsynaptic depolarization causes removal of the Mg^{2+} block (Mayer et al. 1984, Nowak et al. 1984), which together allow Ca^{2+} influx though the NMDA receptors. The level and time course of postsynaptic Ca^{2+} rise depend on the induction protocol: HFS leads to fast, large Ca^{2+} influx, whereas LFS leads to prolonged, modest Ca^{2+} rise (Malenka & Bear 2004, Yang et al. 1999). In the Ca^{2+} hypothesis (Artola & Singer 1993, Lisman 1989, Yang et al. 1999), these two types of Ca^{2+} signals cause the activation of separate molecular pathways. Activation of Ca^{2+}/calmodulin-dependent protein kinase II (CaMKII) by large Ca^{2+} rise is required for LTP, whereas recruitment of phosphatases such as protein phosphatase 1 (PP1) and calcineurin by modest Ca^{2+} increase is necessary for LTD (Malenka & Bear 2004). Spike timing–dependent LTP (tLTP) and LTD (tLTD) also depend on NMDA

STDP: spike timing–dependent plasticity

N-methyl-d-aspartate (NMDA) receptor: subtype of glutamate receptors

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

BAP: back-propagating action potential

AP: action potential

VDCC: voltage-dependent Ca^{2+} channel

receptor activation and the rise in postsynaptic Ca^{2+} level (Bi & Poo 1998, Debanne et al. 1998, Feldman 2000, Magee & Johnston 1997, Markram et al. 1997, Sjostrom et al. 2001, Zhang et al. 1998). Can this simple model for conventional LTP and LTD account for STDP, in particular for its temporal window on the time scale of tens of milliseconds?

tLTP Window

Induction of tLTP requires activation of the presynaptic input milliseconds before the backpropagating action potential (BAP) in the postsynaptic dendrite (pre → post, positive intervals). The BAP can facilitate Mg^{2+} unblocking of NMDA receptors and thus allow Ca^{2+} influx, leading to tLTP induction. However, the width of the tLTP window cannot be explained solely by the time course of NMDA receptor activation. The dissociation of glutamate from the NMDA receptors occurs on the order of hundreds of milliseconds (Lester et al. 1990), much longer than the observed tLTP windows (**Figure 1*a***). The short duration of the window may be due to the kinetics of Mg^{2+} unblocking NMDA receptors (Kampa et al. 2004), such that the BAPs arriving soon after the onset of the excitatory postsynaptic potential (EPSP) are better able to open the NMDA receptors.

In addition to the Mg^{2+} unblock of NMDA receptors, the tLTP window could also be shaped by other types of interactions between the EPSP and the BAP. For example, the EPSP can cause changes in the dendritic conductances that affect the action potential (AP) backpropagation into the dendrites. In the hippocampus, the distal dendrites of CA1 pyramidal neurons express a high density of A-type K^+ channels, which regulate the BAP amplitude (Hoffman et al. 1997). An EPSP that depolarizes the dendrite and inactivates these channels can boost the BAPs arriving within tens of milliseconds (Magee & Johnston 1997, Watanabe et al. 2002). This boosting of the BAPs can in turn increase the Ca^{2+} influx through voltage-dependent Ca^{2+} channels (VDCCs), which can modulate the magnitude of tLTP (Bi & Poo 1998, Froemke et al. 2006, Magee & Johnston 1997). In the neocortex, a similar boosting of the BAP by the preceding EPSP is achieved by voltage-gated Na^{2+} channel activation in the distal dendrites (Stuart & Hausser 2001). Such nonlinear interactions between the EPSP and BAP at short positive intervals could explain the supralinear summation of Ca^{2+} influx to the active synapse in both hippocampal (Magee & Johnston 1997) and neocortical (Koester & Sakmann 1998) (Nevian & Sakmann 2004) neurons.

tLTD Window

Models based on the Ca^{2+} hypothesis have also been used to explain the tLTD window (post → pre, negative intervals) (Karmarkar & Buonomano 2002, Shouval et al. 2002). Assuming that the BAP contains an afterdepolarization lasting for tens of milliseconds and that all relevant Ca^{2+} enters the postsynaptic cell through NMDA receptors, the tLTD window can be explained by the interaction between the EPSP and the BAP. Unlike pairing of the BAP and the EPSP at positive intervals, which causes large Ca^{2+} influx through the NMDA receptors, the EPSP coinciding with the afterdepolarization leads to a moderate Ca^{2+} influx, resulting in tLTD. It should be noted that this model predicts an additional tLTD window at positive intervals outside the tLTP window (**Figure 1*a***, II), where the rise in postsynaptic Ca^{2+} falls within the range for LTD induction. This additional tLTD window has indeed been observed in hippocampal CA1 neurons (Nishiyama et al. 2000, Wittenberg & Wang 2006) but not at other synapses. This suggests a distinct form of STDP at hippocampal synapses, or it could reflect insufficient sampling of long positive intervals in the experimental studies of STDP in other circuits.

In another model for tLTD based on the Ca^{2+} hypothesis (Froemke et al. 2005), a BAP preceding an EPSP induces Ca^{2+} influx through VDCCs, which inactivates the NMDA receptors (Rosenmund et al. 1995, Tong et al. 1995). The reduced Ca^{2+} influx

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

through NMDA receptors in turn leads to tLTD. This model is supported by the observations that tLTD induction requires activation of VDCCs (Bender et al. 2006, Bi & Poo 1998, Froemke et al. 2005, Nevian & Sakmann 2006) and that pairing EPSPs and BAPs at negative intervals leads to sublinear summation of Ca^{2+} influx (Koester & Sakmann 1998, Nevian & Sakmann 2004). Furthermore, in L2/3 pyramidal neurons in visual cortical slices, BAP-induced Ca^{2+}-dependent NMDA receptor inactivation varied with dendritic location, mirroring the location dependence of the tLTD window at these synapses (Froemke et al. 2005).

In some other synapses, tLTD induction does not depend on activation of postsynaptic NMDA receptors (Bender et al. 2006, Egger et al. 1999, Nevian & Sakmann 2006, Sjostrom et al. 2003). These studies suggest a model involving two coincidence detectors, with the NMDA receptor for tLTP and an additional coincidence detector for tLTD. In a two-detector model proposed by Karmarkar & Buonomano (2002), tLTD induction requires activation of postsynaptic mGluRs (metabotropic glutamate receptors) and Ca^{2+} influx through VDCCs, a premise supported by experimental findings in the barrel cortex (Bender et al. 2006, Egger et al. 1999, Nevian & Sakmann 2006). Signaling through mGluRs can lead to phospholipase C (PLC) activation, and Ca^{2+} influx through VDCCs can facilitate mGluR-dependent-PLC activation (Hashimotodani et al. 2005, Maejima et al. 2005). Thus, PLC can serve as a potential coincidence detector for tLTD.

Downstream of coincidence detection, PLC may generate inositol 1,4,5-triphosphate (IP_3), which in turn triggers release of Ca^{2+} from internal stores through IP_3 receptors (IP_3Rs) (Bender et al. 2006). Both PLC activation and Ca^{2+} level elevation (due to influx through VDCCs and/or NMDA receptors, or release from internal stores) can promote endocannabinoid synthesis and release (Hashimotodani et al. 2007). Endocannabinoids play important roles in both short- and long-term depression of many synapses (Chevaleyre et al. 2006). Signaling through presynaptic CB1 endocannabinoid receptors is also required for tLTD for several excitatory–excitatory (Bender et al. 2006; Nevian & Sakmann 2006; Sjostrom et al. 2003) and excitatory–inhibitory connections (Tzounopoulos et al. 2007), presumably by inhibiting presynaptic transmitter release. In **Figure 2**, we have outlined the major signaling pathways implicated in STDP.

mGluR: metabotropic glutamate receptor

STDP OF INHIBITION

Balanced excitation and inhibition are crucial for normal brain functions (Shu et al. 2003) and

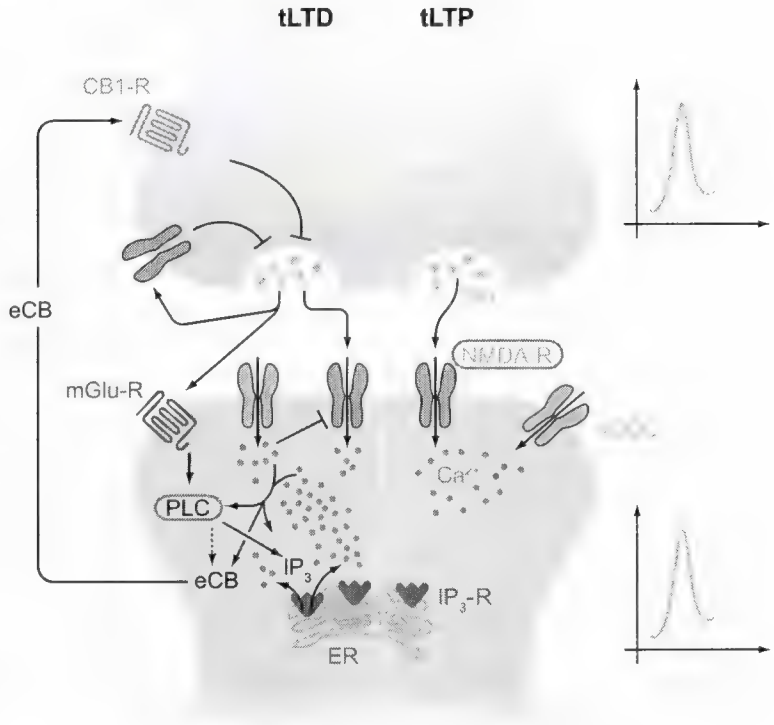

Figure 2

Schematic representation of signaling pathways involved in STDP induction. In tLTP induction (*right*), the NMDA receptors act as coincidence detectors for pre- and postsynaptic spiking. In tLTD induction (*left*) the coincidence detector may vary across synapses. The diagram includes several pathways that have been suggested to play a role in tLTD. Red oval indicates possible coincidence detectors. Arrow indicates activation/potentiation. Blunt-ended line indicates inhibition/suppression. Abbreviations: eCB, endocannabinoids; ER, endoplasmic reticulum; Glu, glutamate; IP_3, inositol 1,4,5-triphosphate; PLC, phospholipase C; VDCCs, voltage-dependent Ca^{2+} channels.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

for regulating experience-dependent developmental plasticity (Hensch 2005). Although the strength of excitatory synapses can be modified through STDP, an important question is whether and how correlated pre- and postsynaptic activity affects inhibitory circuits. Inhibition in a network depends on both the excitatory synapses onto inhibitory neurons and the inhibitory synapses themselves. Spike timing–dependent plasticity has been studied at both of these synapses.

STDP of Excitatory Synapses onto Inhibitory Neurons

In a cerebellum-like structure in the electric fish, Bell and colleagues (1997) measured the excitatory inputs to Purkinje-like GABAergic neurons to study the dependence of synaptic modification on the temporal order of pre- and postsynaptic spiking. Pre → post pairing within a 60-ms window induces LTD, whereas post → pre pairing leads to LTP (**Figure 1*b***, I). This asymmetrical window is thus opposite in polarity to the STDP window for the synapses between excitatory neurons (**Figure 1*a***, I). However, given the difference in the postsynaptic neurons, the functional consequences of the two learning rules may be similar, and they could act cooperatively in activity-dependent network modifications. Mechanistically, LTD induced by pre → post pairings required NMDA receptor activation and postsynaptic Ca^{2+} elevation (Han et al. 2000), similarly to tLTD for excitatory-excitatory connections. However, LTP of these synapses can be induced by EPSPs alone without postsynaptic spiking, indicating a nonassociative component of the synaptic plasticity.

Another study of excitatory inputs to inhibitory neurons was conducted in mouse brain stem slices by pairing parallel fiber stimulation with cartwheel neuron spiking (Tzounopoulos et al. 2004, 2007). Pre → post pairing within a narrow window (<10 ms) induces LTD, whereas post → pre pairing causes no change in synaptic strength (**Figure 1*b***, II). For these synapses, LTD depends on postsynaptic NMDA receptor activation, Ca^{2+} influx, and endocannabinoid signaling, similar to the findings at excitatory synapses onto pyramidal neurons (see previous section). Interestingly, synapses from the same presynaptic fibers onto excitatory postsynaptic neurons (fusiform principal neurons) exhibit a STDP window similar to that of other excitatory-excitatory connections (**Figure 1*a***, I) (Tzounopoulos et al. 2004, 2007). This target specificity of the learning rule can be attributed to the selective distribution of presynaptic endocannabinoid CB1 receptors in different axonal terminals.

STDP of GABAergic Synapses

Compared with the glutamatergic synapses, the learning rules for GABAergic synapses appear more variable. In a study of inhibitory inputs to neocortical L2/3 pyramidal neurons, synaptic modification was induced by pairing single presynaptic spikes with high-frequency postsynaptic bursts. Overlapping pre- and postsynaptic spiking induced LTD, and nonoverlapping post → pre spiking within hundreds of milliseconds induced LTP (Holmgren & Zilberter 2001) (**Figure 1*c***, I). In the hippocampus, GABAergic synapses onto CA1 pyramidal neurons exhibit a symmetrical window, with pairing of single pre- and postsynaptic spikes at short intervals (within ±20 ms) leading to LTP, and pairing at long intervals leading to LTD (Woodin et al. 2003) (**Figure 1*c***, II). In contrast, in the entorhinal cortex GABAergic inputs to layer II excitatory stellate cells exhibit an asymmetric window similar to the STDP window for excitatory-excitatory connections: LTP was found at positive intervals and LTD at negative intervals (Haas et al. 2006) (**Figure 1*c***, III). Despite the differences between these temporal windows for GABAergic synapses, both the induction mechanism and the loci of expression have similarities. In both hippocampal CA1 (Woodin et al. 2003) and the entorhinal cortex (Haas et al. 2006), the induction of synaptic modification

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

depends on postsynaptic Ca^{2+} influx through the L-type Ca^{2+} channels, and presynaptic expression was excluded because no change was observed in the paired pulse ratio. In the hippocampus (Woodin et al. 2003), the changes in inhibitory postsynaptic current (IPSC) amplitude are due to changes in the Cl^- reversal potential mediated by modification of the KCC2 K^+-Cl^- cotransporter, further indicating that the expression is postsynaptic.

STDP WITH COMPLEX SPIKE PATTERNS

To study synaptic plasticity, the induction paradigms are often selected for their effectiveness rather than for their physiological relevance, thus providing limited information on how circuits are modified by natural patterns of activity. Although most induction protocols for STDP consisted of repetitive pairing of pre- and postsynaptic spikes at regular intervals, neuronal activity in vivo is far from regular (Softky & Koch 1993), with periods of almost no activity intermingled with short bouts of high-frequency spike bursts. During each presynaptic burst, transmitter release is likely to be affected by short-term plasticity (Zucker & Regehr 2002), and in each postsynaptic burst the efficacy of individual spike propagation may depend on the spike pattern (Spruston et al. 1995; Williams & Stuart 2000). How well does the STDP learning rule measured with simple spike patterns account for the synaptic changes induced by naturalistic spike trains? When multiple spike pairs fall within the STDP window, how are the contributions of individual spikes integrated?

One simple strategy to study the interaction among multiple spikes is to add one spike at a time to the existing pairing protocol. In L2/3 of visual cortical slices (Froemke & Dan 2002) and in hippocampal cultures (Wang et al. 2005), spike "triplets" (pre → post → pre or post → pre → post) and "quadruplets" (pre → post → post → pre or post → pre → pre → post) were used to induce synaptic modifications. In both studies, the interaction between multiple spikes was nonlinear, but the specific forms of nonlinearity were different. In cortical L2/3, the nonlinear interactions could be accounted for by a suppression model, in which the efficacy of later spikes in each train for synaptic modification is reduced by the preceding spikes (Froemke & Dan 2002). This model accurately predicted the synaptic changes induced by natural spike trains recorded in vivo in response to visual stimulation. In cultured hippocampal neurons, the "pre → post → pre" triplets induce no synaptic change, which suggests that LTP and LTD cancel each other, but the "post → pre → post" triplets induce LTP, which suggests that LTP "wins over" LTD under this condition. A third study using spike triplets showed that in hippocampal slices, different learning rules are revealed with different numbers of spike pairings (Wittenberg & Wang 2006). With 20–30 pairings at 5 Hz, LTP was induced regardless of the temporal order of the spikes. With 70–100 repeats, however, LTP was observed at short positive intervals (<30 ms), and LTD was found at both negative intervals and at long positive intervals (>30 ms) (**Figure 1*a***, II). These results suggest that the integration across multiple spike pairs depends on the activity patterns over several minutes.

The effects of pre- and/or postsynaptic spike bursts on synaptic modification have also been examined. Paired recordings from L5 pyramidal neurons in visual cortical slices showed that the synaptic change depends on both the spike frequency within each burst and the interval between the pre- and postsynaptic spikes (Sjostrom et al. 2001). At high frequencies (≥50 Hz), LTP is induced regardless of the pre/post interval, whereas at intermediate frequencies (10–40 Hz), the pre/post interval determines the sign and magnitude of synaptic modification as described by the STDP window (**Figure 1*a***, I). Pairing at low frequencies (<1 Hz) notably fails to induce LTP. This is likely caused by the small EPSPs evoked by activating a single presynaptic neuron in paired recordings because LTP can be rescued by adding extracellular stimulation that provides additional

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

depolarization. The combined dependence of synaptic modification on burst timing and frequency can be accounted for by a model in which LTP wins over LTD, and only the interactions between neighboring spikes contribute to synaptic modification (Sjostrom et al. 2001). In another study in L2/3 neurons in rat visual cortical slices (Froemke et al. 2006), pairing of pre- and postsynaptic bursts at high frequencies also favored LTP regardless of the pre/post spike timing. However, systematic examination of the dependence of synaptic modification on both the number and the timing of pre- and postsynaptic spikes led to a modified suppression model (Froemke et al. 2006), which incorporates short-term depression of the presynaptic input (Zucker & Regehr 2002) and frequency-dependent attenuation of postsynaptic spikes (Spruston et al. 1995). Note that in both models described above, burst-induced synaptic modification is accounted for by integrating the contributions of individual spike pairs. However, in some synapses the learning rule for bursts seems to be completely different from that for individual spikes (Birtoli & Ulrich 2004, Kampa et al. 2006, Pike et al. 1999).

Although the above studies focused on synaptic modifications induced by short bursts lasting for tens of milliseconds, in some circuits bursts can last for hundreds of milliseconds to several seconds. In the developing retinogeniculate synapse, bursts of retinal ganglion cells lasting seconds are believed to be critical for circuit refinement (Butts & Rokhsar 2001). Temporally overlapping pre- and postsynaptic bursts (interval within a window of ~1 s) result in synaptic potentiation, whereas nonoverlapping bursts cause a slight depression (Butts et al. 2007). The degree of potentiation can be predicted by a model in which LTP depends on the interval but not the order between the pre- and postsynaptic bursts, and it increases linearly with the number of spikes in the burst. This is reminiscent of the classic correlation-based learning rule for synaptic plasticity (Stent 1973). A strikingly similar window for burst timing was found in the hippocampal CA3 region for correlated activation of the associational/commissural (A/C) fibers and the mossy fibers (Kobayashi & Poo 2004), although no depression was observed. In both studies, the width of the temporal window seems to scale with the duration of the spike bursts used in the induction protocol, and the changes in synaptic strength depend on the interburst interval rather than the precise timing of individual spikes. Such burst timing–dependent plasticity rules may be functionally advantageous for the circuits in which the information relevant for synaptic refinement is contained in the timing of the bursts rather than that of individual spikes (Butts & Rokhsar 2001).

Together, the studies described above indicate that the integration across multiple spike pairs for the induction of synaptic modification is highly nonlinear. The nature of the nonlinear interaction is likely to depend on short-term plasticity of the presynaptic neurons, on the biophysical properties of the postsynaptic dendrites, and on the downstream signaling pathways present in different cell types. Further characterization of the diversity of integration mechanisms for STDP will allow better understanding of circuit remodeling induced by natural patterns of neuronal activity.

DEPENDENCE ON DENDRITIC LOCATION

In the central nervous system, each neuron may receive thousands of synaptic inputs distributed throughout its dendritic tree. The processing of each input depends on the dendritic location (Hausser & Mel 2003) owing to both the passive cable properties (Rall 1967) and the nonuniform distribution of active conductances (Migliore & Shepherd 2002). Such location-dependent processing and integration of synaptic inputs are believed to be essential aspects of neuronal computation. Since a hallmark of STDP is its dependence on the BAPs, which are strongly attenuated along the dendrite (Stuart & Sakmann 1994, Stuart et al. 1997b, Waters et al. 2005), synaptic modification is likely to vary with

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

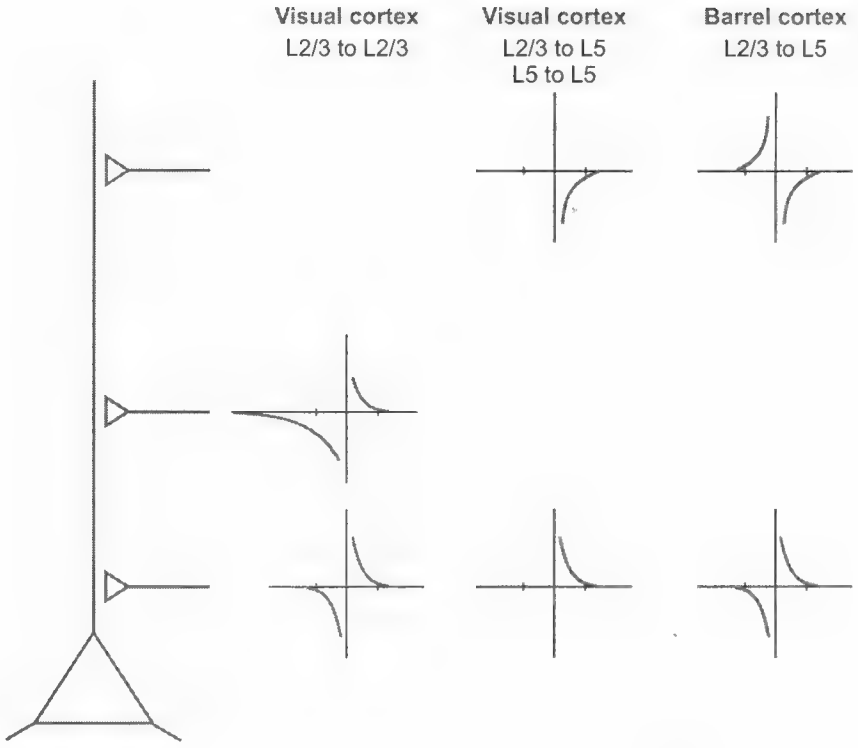

Figure 3
Dependence of STDP on dendritic location.

dendritic location (Rao & Sejnowski 2001b). Recent studies have examined the location dependence of both tLTP and tLTD. In L2/3 of rat visual cortex, the magnitude of tLTP induced by pre → post pairing of single spikes was smaller at intermediate-distal (100–150 μm) than at proximal (<50 μm) segments of the apical dendrite (Froemke et al. 2005) (**Figure 3**, left column). This reduction of tLTP amplitude is likely due to distance-dependent attenuation of the BAP. In experiments with paired recordings from a L5 and a L2/3 pyramidal neuron or from two L5 neurons (Sjostrom & Hausser 2006), burst pairing at positive intervals led to LTP at the proximal synapses but LTD at the distal synapses (**Figure 3**, middle column). Similar location dependence was also found among L2/3 to L5 connections by pairing a single EPSP with a postsynaptic burst at positive intervals (Letzkus et al. 2006) (**Figure 3**, right column). BAP boosting by subthreshold local dendritic depolarization or extracellular stimulation recovered tLTP at distal synapses (Letzkus et al. 2006, Sjostrom & Hausser 2006), which suggests that distal tLTP requires cooperativity among inputs.

Two distinct effects have been reported for post → pre pairing. In L2/3 pyramidal neurons, the width of the tLTD window measured with single spike pairing is broader for intermediate-distal than for proximal inputs (Froemke et al. 2005) (**Figure 3**, left column). This difference in width is correlated with the window for AP-induced suppression of NMDA receptor activation, which suggests that the suppression plays an important role in setting the tLTD window. In L2/3–L5 synapses in rat barrel cortex, pairing single presynaptic spikes with postsynaptic bursts at negative intervals leads to LTD at proximal locations but LTP of distal inputs (Letzkus et al. 2006) (**Figure 3**, right column). This distal LTP could be explained by the induction of dendritic Ca^{2+} spikes by the later BAPs in the burst (Larkum et al. 1999a, Stuart et al. 1997a), such that the EPSP coincides with the peak postsynaptic depolarization. Local dendritic spikes can also play a prominent role in coincidence detection in the neocortex

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

(Larkum et al. 1999b) and in LTP induction in hippocampal CA1 (Golding et al. 2002) and the amygdala (Humeau & Luthi 2007).

Comparison across these studies suggests that the degree of spatial variation of the learning rule depends on the dendritic morphology, with quantitative changes over short distances (e.g., dendrite of L2/3 neurons) and qualitative differences along long dendrites (e.g., apical dendrites of L5 pyramids). Although the dendritic variations of STDP summarized above can be explained largely by differences in the local active conductances, the backpropagation of APs, or the local generation of Ca^{2+} spikes, differential distribution of other pre- and postsynaptic molecular machineries could also contribute to the observed heterogeneity. Functionally, the spatial variation of the STDP rule may lead to differential input selection at distal and proximal dendrites. For example, the relative paucity of LTP at distal dendrites after pre → post pairing predicts that proximal inputs should be stronger than distal inputs (Sjostrom & Hausser 2006). The involvement of locally generated Ca^{2+} spikes in LTP induction (Golding et al. 2002, Kampa et al. 2006) likely rewards cooperativity among distal inputs because their synchronous activation is known to evoke dendritic spikes. Furthermore, the broader LTD window for intermediate distal inputs to L2/3 neurons suggests that the distal dendrites strongly favor transient over prolonged inputs (Froemke et al. 2005).

MODULATION OF STDP BY OTHER INPUTS

In addition to the spiking of the pre- and postsynaptic neurons, STDP is also regulated by other inputs. In particular, neuromodulators and inhibitory activity in the network can affect both the magnitude and the temporal window of STDP.

Neuromodulators such as norepinephrine and acetylcholine (ACh) play important roles in experience-dependent neural plasticity (Bear & Singer 1986, Kilgard & Merzenich 1998). At the cellular level, neuromodulators can influence AP backpropagation by modulating the activation and inactivation of various active conductances (Johnston et al. 1999). For example, agonists to muscarinic ACh receptors can reduce spike attenuation during high-frequency bursts, probably through reduction of Na^+ channel inactivation (Johnston et al. 1999, Tsubokawa & Ross 1997). Both β-adrenergic and muscarinic ACh receptor agonists can boost AP backpropagation by downregulating transient K^+ channels through protein kinase A (PKA) and protein kinase C (PKC) activation, respectively (Hoffman & Johnston 1998, 1999). Dopamine also has a similar effect on the BAP (Hoffman & Johnston 1999).

Such modulations of the BAPs are likely to have profound effects on STDP, particularly at distal dendritic locations. In the Schaffer collateral pathway to hippocampal CA1, pairing a weak and a strong input (which evokes postsynaptic spiking) at positive intervals can induce NMDA receptor–dependent tLTP within a narrow window of 3–10 ms. Bath application of isoproterenol, a β-adrenergic receptor agonist, broadens the window to 15 ms without changing the magnitude of tLTP (Lin et al. 2003), an effect that depends on PKA and mitogen-activated protein kinase (MAPK) signaling. In the amygdala, dopamine can gate the induction of tLTP by suppressing feedforward inhibitory inputs to the postsynaptic cell (Bissiere et al. 2003). In L5 pyramidal neurons of the prefrontal cortex, nicotine application converted tLTP to tLTD by reducing dendritic Ca^{2+} signals during spike pairing (Couey et al. 2007), and this reduction is mediated by an enhancement of GABAergic synaptic transmission. In L2/3 pyramidal neurons, activation of M1 muscarinic receptors promotes tLTD induction through a PLC-dependent pathway, whereas β-adrenergic receptor activation promotes tLTP through the adenylate cyclase cascade (Seol et al. 2007). Thus, neuromodulators can regulate both the magnitude and the polarity of synaptic modifications.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

The timing and location of inhibitory inputs can also affect STDP. Somatic inhibition can prevent AP propagation through hyperpolarization and shunting (Miles et al. 1996, Tsubokawa & Ross 1996), which may preclude STDP induction. In contrast, inhibitory inputs to the dendrites have a variety of effects, from reducing dendritic depolarization through shunting to facilitating depolarization and even spike generation (Gulledge & Stuart 2003). An additional layer of complexity is added by the fact that the strength and distribution of inhibition are developmentally regulated (Hensch 2005), predicting that the learning rule can vary considerably across developmental stages. In hippocampal CA1 pyramidal neurons, pairing single pre- and postsynaptic spikes at positive intervals leads to tLTP in juvenile (p9–p14) but not in young (p22–p28) rats (Meredith et al. 2003). However, in young rats tLTP can be rescued by replacing the single postsynaptic spike with a burst or by adding $GABA_A$ antagonists, suggesting that the change in tLTP threshold might be due to a developmental enhancement of inhibition in this circuit.

PLASTICITY OF NEURONAL EXCITABILITY AND SYNAPTIC INTEGRATION

Information processing by neuronal networks depends not only on the connectivity between neurons, but also on the intrinsic conductances in each neuron that determine its excitability and synaptic integration. Changes in neuronal excitability have been reported in a variety of invertebrate and vertebrate neural circuits during associative learning (Daoudal & Debanne 2003, Zhang & Linden 2003). At the cellular level, LTP induction by tetanic stimulation also leads to increases in intrinsic excitability in both the hippocampus and the cerebellum (Aizenman & Linden 2000, Armano et al. 2000, Bliss & Gardner-Medwin 1973, Bliss & Lomo 1973). These activity-dependent changes in intrinsic neuronal properties may interact synergistically with synaptic plasticity to mediate learning and memory.

Changes in neuronal excitability have also been examined in the context of STDP. In hippocampal cell cultures (Ganguly et al. 2000) and neocortical slices (Li et al. 2004), repeated pre → post pairing of single spikes leads to LTP and to an enhancement of excitability and spike time reliability of the presynaptic neurons. Pairings at negative intervals result in LTD and a reduction in presynaptic excitability (Li et al. 2004). Mechanistically, these presynaptic changes require NMDA receptor activation and Ca^{2+} influx to the postsynaptic neuron, suggesting the involvement of retrograde signaling. On the presynaptic side, PKC is necessary for the increase in excitability (Ganguly et al. 2000), and both PKC and PKA are required for the decrease (Li et al. 2004). Interestingly, the changes in excitability can be dissociated from the changes in synaptic strength because presynaptic blockage of PKC and/or PKA abolished the excitability changes with little effect on the synaptic modifications.

Activity-dependent changes in intrinsic membrane properties can also affect synaptic integration (Magee & Johnston 2005). A recent study examined the changes in spatial summation between two input pathways in hippocampal CA1 neurons following STDP induction (Wang et al. 2003). Induction of tLTP in one pathway resulted in an increase in the linearity of spatial summation of the two pathways, whereas induction of tLTD produced the opposite effect. The observed changes depend on NMDA receptor activation and may be mediated by modifications of the I_h channels. In another study in hippocampal CA1, LTP induction by paired theta bursts causes an increase in the linearity of temporal summation between the potentiated input and a neighboring input (Xu et al. 2006); the temporal specificity of this effect varied with dendritic location. For distal inputs, the increase in linearity is limited to EPSPs arriving within 5 ms of each other, favoring summation of coincident inputs. In contrast, for proximal inputs the increase can be observed for EPSPs arriving within

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

a broader window of 20 ms. Such location-dependent modulation of synaptic integration may interact with the location dependence of the STDP learning rule (see above) to further enrich dendritic processing.

STDP IN VIVO

Whereas most of the early experiments on STDP were conducted in slices and cell cultures, an increasing number of studies have begun to address the functional consequences of STDP in intact nervous systems. Neural circuits in vivo exhibit both spontaneous activity and sensory-evoked responses, modulated by the behavioral states of the animal. Backpropagation of the APs may also be more variable in vivo, as the neurons receive barrages of excitatory and inhibitory inputs (Destexhe et al. 2003). These factors could significantly complicate the rules for synaptic plasticity. How well does the STDP learning rule described in vitro apply to activity-dependent synaptic modification in vivo?

Electrical Stimulation

The first demonstration of STDP in vivo came from a study at the retinotectal projection in the developing *Xenopus* (Zhang et al. 1998). Repetitive electrical stimulation of the retinal ganglion cells within 20 ms before tectal neuron spiking leads to LTP, whereas pairings at negative intervals lead to LTD. Both LTP and LTD are NMDA receptor dependent, and the temporal window is similar to the STDP windows measured in vitro (e.g., Bi & Poo 1998, Froemke & Dan 2002, Tzounopoulos et al. 2004). In addition to the strength of the retinotectal connection, the amplitude of the tectal visual response can also be modified by pairing visual stimulation with postsynaptic spiking (Mu & Poo 2006, Vislay-Meltzer et al. 2006).

Plasticity with similar asymmetric windows has also been demonstrated in the mammalian visual cortex. Optical imaging in the kitten visual cortex showed that pairing visual stimulation at a given orientation with cortical electrical stimulation leads to changes in the orientation map (Schuett et al. 2001). Electrical activation after the arrival of the visual input causes expansion of the cortical representation of the paired orientation, whereas the reverse order causes a reduction. Whole-cell recordings in juvenile rat visual cortex showed that pairing visual stimulation with single neuron spiking leads to potentiation or depression of the visual response, depending on the order between the visual inputs and the postsynaptic spiking (Meliza & Dan 2006).

STDP has also been described in other sensory modalities in vivo. In the somatosensory cortex of anesthetized rats, pairing subthreshold whisker deflections with postsynaptic spiking at negative intervals leads to LTD of the paired whisker (Jacob et al. 2007). In an olfactory circuit of the locusts (β-lobe in the mushroom body), pairing odor-induced synaptic activity with postsynaptic spiking results in robust synaptic modifications, with a temporal window similar to those for vertebrate excitatory synapses (**Figure 1*a***, I) (Cassenaer & Laurent 2007).

In the motor system, STDP has been demonstrated in human subjects. Pairing electrical stimulation of a somatosensory afferent nerve with transcranial magnetic stimulation (TMS) of the motor cortex leads to long-lasting changes in the motor-evoked potentials (MEPs) elicited by TMS (Wolters et al. 2003). The direction and magnitude of the change depend on the relative timing between the afferent stimulation and the TMS within a window of tens of milliseconds, comparable to the STDP windows measured in vitro. The potentiation induced by pairing at positive intervals can be blocked by NMDA receptor antagonists (Stefan et al. 2002), and the depression at negative intervals is blocked by both NMDA receptor and VDCC antagonists (Wolters et al. 2003), consistent with the pharmacological properties of STDP found in several studies (Bi & Poo 2001). Wolters et al. (2005) also used a similar experimental protocol to demonstrate STDP in human somatosensory cortex.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Paired Sensory Stimulation

Although electrical stimulation affords excellent control of spike timing in the study of STDP, an important question is whether the temporal requirements of this learning rule can be satisfied under natural conditions, as spiking responses to sensory stimuli are known to be highly variable (Shadlen & Newsome 1994). Several studies on the functional role of STDP in vivo have been performed with pure sensory stimulation. In anesthetized adult cats, repetitive presentation of gratings at a pair of orientations induced shifts in orientation tuning of individual V1 neurons; the direction of the shift depended on the temporal order of the two orientations (Yao et al. 2004, Yao & Dan 2001). In a parallel set of experiments in the space domain, repeated visual stimulation in two adjacent retinal regions induced shifts in V1 receptive fields (Fu et al. 2002), with a similar dependence on the stimulus order. In both the orientation and space domain, significant changes in cortical response properties were observed at intervals within ±40 ms, similar to the STDP windows observed in vitro. For the shift in orientation tuning, the effect showed complete interocular transfer, indicating that the underlying neuronal modifications occur largely in the cortex, after the inputs from the two eyes converge (Yao et al. 2004). Psychophysical experiments in human subjects using analogous induction protocols showed perceptual changes consistent with the electrophysiological effects (Fu et al. 2002, Yao et al. 2004, Yao & Dan 2001), which suggests that the neuronal changes have direct consequences in visual perception.

Motion Stimuli

Compared with the repetitively flashed stimuli used in the above studies, moving stimuli are much more common in nature. Motion stimuli are intrinsically sequential (e.g., an object moving across the visual field should sequentially enter the neuronal receptive fields distributed along its trajectory) and are thus ideally suited for interacting with the STDP learning rule. In the *Xenopus* tadpole, repeated presentation of a moving bar in a given direction selectively potentiated the response to the conditioned direction, resulting in the emergence of direction sensitivity in the tectal neurons (Engert et al. 2002). Induction of direction selectivity through STDP has indeed been predicted in a theoretical study (Rao & Sejnowski 2001a). A follow-up experiment using both sequentially flashed bars and moving bars provided further support for the role of STDP in the induction of direction selectivity (Mu & Poo 2006). The selective enhancement at the conditioned direction manifests as a potentiation of the early phase and a reduction of the late phase of the visual response, consistent with the prediction from STDP. Blocking the cellular signaling pathways underlying STDP abolished the effect of unidirectional motion stimuli in inducing direction selectivity.

In the visual cortex, the interaction between motion stimuli and STDP has been used to predict two receptive field properties and to explain two motion-position illusions. Model simulations predicted that the prevalence of motion stimuli in various directions during visual cortical development would lead to a spatial asymmetry in the direction-selective inputs to each cortical neuron (e.g., inputs preferring rightward motion are biased toward the left side of the receptive field) (Fu et al. 2004). This asymmetry in the mature cortex in turn predicts that (*a*) receptive field position depends on the local motion signals within the test stimuli, and (*b*) motion adaptation causes the receptive field position to shift. Both effects were confirmed experimentally in anesthetized cat V1. Psychophysical measurement using matching stimulus parameters showed that these physiological effects could each explain a known visual illusion involving the interaction between motion and perceived object position (De Valois & De Valois 1991, Nishida & Johnston 1999, Ramachandran & Anstis 1990, Snowden 1998, Whitaker et al. 1999).

In addition to the motion signals in sensory inputs, locomotion of the animal may also induce circuit modification through STDP. The

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

place fields of hippocampal neurons are known to be dynamically modified as the animal navigates in a novel environment. During repeated running of a linear track, the place fields of both CA1 and CA3 cells are initially symmetrical, but they experience a gradual asymmetric expansion against the direction of locomotion (Lee et al. 2004; Mehta et al. 1997, 2000). Simulation with a simple feedforward network model showed that this effect can be explained by STDP (Blum & Abbott 1996, Mehta et al. 2000). In the orientation domain, Yu et al. (2006) recently reported a similar shift in head-direction tuning curves in thalamic head-direction cells as the animal runs in a circular track.

Sensory Deprivation

STDP may also play a role in other forms of experience-dependent plasticity, even if the sensory inputs do not explicitly involve timing on the order of tens of milliseconds. In an experiment measuring the neural activity during sensory deprivation, rats were chronically implanted with electrode arrays to monitor the spiking activity in L4 and L2/3 of the barrel cortex during free-moving behaviors (Celikel et al. 2004). Stimulus deprivation induced by trimming a single whisker, a manipulation known to induce whisker map reorganization, caused an immediate reversal of the firing order and decreased correlation between L4 and L2/3 neurons. Both of these changes are known to drive tLTD in barrel cortical slices (Feldman 2000), thus providing a plausible explanation for deprivation-induced LTD of L4 to L2/3 connections (Allen et al. 2003). In addition to the somatosensory system, sensory deprivation induces circuit reorganization in the visual and auditory systems (Buonomano & Merzenich 1998, Gilbert 1998). It would be interesting to test whether deprivation in these modalities (e.g., monocular deprivation of visual input) also induces changes in the relative spike timing among neurons that could cause the observed circuit modifications through STDP.

FINAL REMARKS

Over the past decade, the STDP learning rule has been demonstrated in a range of species from insects to humans, and our understanding of its cellular mechanisms and functional implications has progressed significantly. However, many questions remain unresolved.

Regarding the mechanism, it remains unclear whether a single model can explain STDP at different synapses or whether different neurons employ distinct molecular machineries to achieve similar outcomes. Studies are only beginning to examine whether and how STDP depends on several signaling events that have been strongly implicated in conventional LTP and LTD, including secretion of brain-derived neurotrophic factor (BDNF) and nitric oxide (Mu & Poo 2006), activation of CaMKII (Tzounopoulos et al. 2007) and phosphatases (Froemke et al. 2005), and modification and insertion/removal of AMPA receptors. It would also be interesting to investigate whether the type of NMDA receptor subunits (NR2A/NR2B) and their synaptic location play a role in STDP (Sjostrom et al. 2003), as has been suggested for LTP/LTD induced by HFS/LFS (Cull-Candy & Leszkiewicz 2004, Liu et al. 2004). In addition, whereas several molecules have been proposed as coincidence detectors at excitatory synapses (**Figure 2**), there is so far no candidate for inhibitory synapses. Furthermore, although postsynaptic Ca^{2+} signals are required for STDP in most cell types, recent imaging experiments showed that volume-averaged Ca^{2+} transients in the dendritic spines are poorly correlated with the direction of synaptic modification (Nevian & Sakmann 2006). Perhaps new techniques that allow measurement of Ca^{2+} signals at a more microscopic scale (e.g., microdomains) will shed new light on the cellular mechanisms of STDP.

To understand the functional consequences of STDP, an important factor to consider is the high level of ongoing activity in vivo. Spontaneous activity can significantly affect membrane potential, conductance, and intracellular

Ca^{2+} levels, and in some cases it can boost AP backpropagation in vivo (Waters & Helmchen 2004). These effects will likely modulate the rules for synaptic modification. Furthermore, spontaneous postsynaptic spiking reduces the persistence of synaptic potentiation and depression (Zhou et al. 2003). An important question is how experience-dependent synaptic modifications can persist in vivo in the face of the ongoing network activity. Recent studies have suggested that sensory-evoked activity patterns can reverberate in subsequent spontaneous activity in early sensory circuits (Galan et al. 2006, Yao et al. 2007) or be replayed in the hippocampus during sleep (Ji & Wilson 2007, Louie & Wilson 2001, Nadasdy et al. 1999, Ribeiro et al. 2004, Wilson & McNaughton 1994). These reactivated patterns may serve to consolidate the transient effects of sensory stimulation into long-lasting circuit modifications. Characterization of neuronal plasticity at the network level during natural behaviors is a crucial step in understanding the neural basis for learning and memory.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

DISCLOSURE STATEMENT

The authors are not aware of any biases that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

Abbott LF, Nelson SB. 2000. Synaptic plasticity: taming the beast. *Nat. Neurosci.* 3(Suppl):1178–83

Aizenman CD, Linden DJ. 2000. Rapid, synaptically driven increases in the intrinsic excitability of cerebellar deep nuclear neurons. *Nat. Neurosci.* 3:109–11

Allen CB, Celikel T, Feldman DE. 2003. Long-term depression induced by sensory deprivation during cortical map plasticity in vivo. *Nat. Neurosci.* 6:291–99

Armano S, Rossi P, Taglietti V, D'Angelo E. 2000. Long-term potentiation of intrinsic excitability at the mossy fiber-granule cell synapse of rat cerebellum. *J. Neurosci.* 20:5208–16

Artola A, Brocher S, Singer W. 1990. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* 347:69–72

Artola A, Singer W. 1987. Long-term potentiation and NMDA receptors in rat visual cortex. *Nature* 330:649–52

Artola A, Singer W. 1993. Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends Neurosci.* 16:480–87

Bear MF, Singer W. 1986. Modulation of visual cortical plasticity by acetylcholine and noradrenaline. *Nature* 320:172–76

Bell CC, Han VZ, Sugawara Y, Grant K. 1997. Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387:278–81

Bender VA, Bender KJ, Brasier DJ, Feldman DE. 2006. Two coincidence detectors for spike timing-dependent plasticity in somatosensory cortex. *J. Neurosci.* 26:4166–77

Bi G, Poo M. 2001. Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu. Rev. Neurosci.* 24:139–66

Bi GQ, Poo MM. 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* 18:10464–72

Birtoli B, Ulrich D. 2004. Firing mode-dependent synaptic plasticity in rat neocortical pyramidal neurons. *J. Neurosci.* 24:4935–40

Bissiere S, Humeau Y, Luthi A. 2003. Dopamine gates LTP induction in lateral amygdala by suppressing feedforward inhibition. *Nat. Neurosci.* 6:587–92

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Bliss TV, Collingridge GL. 1993. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361:31–39

Bliss TV, Gardner-Medwin AR. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232:357–74

Bliss TV, Lomo T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* 232:331–56

Blum KI, Abbott LF. 1996. A model of spatial map formation in the hippocampus of the rat. *Neural. Comput.* 8:85–93

Boettiger CA, Doupe AJ. 2001. Developmentally restricted synaptic plasticity in a songbird nucleus required for song learning. *Neuron* 31:809–18

Buonomano DV, Merzenich MM. 1998. Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* 21:149–86

Butts DA, Kanold PO, Shatz CJ. 2007. A burst-based "Hebbian" learning rule at retinogeniculate synapses links retinal waves to activity-dependent refinement. *PLoS Biol.* 5:e61

Butts DA, Rokhsar DS. 2001. The information content of spontaneous retinal waves. *J. Neurosci.* 21:961–73

Cassenaer S, Laurent G. 2007. Hebbian STDP in mushroom bodies facilitates the synchronous flow of olfactory information in locusts. *Nature* 448:709–13

Celikel T, Szostak VA, Feldman DE. 2004. Modulation of spike timing by sensory deprivation during induction of cortical map plasticity. *Nat. Neurosci.* 7:534–41

Chapman PF, Kairiss EW, Keenan CL, Brown TH. 1990. Long-term synaptic potentiation in the amygdala. *Synapse* 6:271–78

Chevaleyre V, Takahashi KA, Castillo PE. 2006. Endocannabinoid-mediated synaptic plasticity in the CNS. *Annu. Rev. Neurosci.* 29:37–76

Clugnet MC, LeDoux JE. 1990. Synaptic plasticity in fear conditioning circuits: induction of LTP in the lateral nucleus of the amygdala by stimulation of the medial geniculate body. *J. Neurosci.* 10:2818–24

Couey JJ, Meredith RM, Spijker S, Poorthuis RB, Smit AB, et al. 2007. Distributed network actions by nicotine increase the threshold for spike-timing-dependent plasticity in prefrontal cortex. *Neuron.* 54:73–87

Cull-Candy SG, Leszkiewicz DN. 2004. Role of distinct NMDA receptor subtypes at central synapses. *Sci. STKE* 2004:re16

Daoudal G, Debanne D. 2003. Long-term plasticity of intrinsic excitability: learning rules and mechanisms. *Learn. Mem.* 10:456–65

Debanne D, Gahwiler BH, Thompson SM. 1998. Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J. Physiol.* 507(Pt. 1):237–47

Destexhe A, Rudolph M, Pare D. 2003. The high-conductance state of neocortical neurons in vivo. *Nat. Rev. Neurosci.* 4:739–51

De Valois RL, De Valois KK. 1991. Vernier acuity with stationary moving Gabors. *Vision Res.* 31:1619–26

Dudek SM, Bear MF. 1993. Bidirectional long-term modification of synaptic effectiveness in the adult and immature hippocampus. *J. Neurosci.* 13:2910–18

Egger V, Feldmeyer D, Sakmann B. 1999. Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat. Neurosci.* 2:1098–105

Engert F, Tao HW, Zhang LI, Poo MM. 2002. Moving visual stimuli rapidly induce direction sensitivity of developing tectal neurons. *Nature* 419:470–75

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Feldman DE. 2000. Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27:45–56

Froemke RC, Dan Y. 2002. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416:433–38

Froemke RC, Poo MM, Dan Y. 2005. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434:221–25

Froemke RC, Tsay IA, Raad M, Long JD, Dan Y. 2006. Contribution of individual spikes in burst-induced long-term synaptic modification. *J. Neurophysiol.* 95:1620–29

Fu YX, Djupsund K, Gao H, Hayden B, Shen K, Dan Y. 2002. Temporal specificity in the cortical plasticity of visual space representation. *Science* 296:1999–2003

Fu YX, Shen Y, Gao H, Dan Y. 2004. Asymmetry in visual cortical circuits underlying motion-induced perceptual mislocalization. *J. Neurosci.* 24:2165–71

Galan RF, Weidert M, Menzel R, Herz AV, Galizia CG. 2006. Sensory memory for odors is encoded in spontaneous correlated activity between olfactory glomeruli. *Neural. Comput.* 18:10–25

Ganguly K, Kiss L, Poo M. 2000. Enhancement of presynaptic neuronal excitability by correlated presynaptic and postsynaptic spiking. *Nat. Neurosci.* 3:1018–26

Gilbert CD. 1998. Adult cortical dynamics. *Physiol. Rev.* 78:467–85

Golding NL, Staff NP, Spruston N. 2002. Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* 418:326–31

Gulledge AT, Stuart GJ. 2003. Excitatory actions of GABA in the cortex. *Neuron* 37:299–309

Haas JS, Nowotny T, Abarbanel HD. 2006. Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J. Neurophysiol.* 96:3305–13

Han VZ, Grant K, Bell CC. 2000. Reversible associative depression and nonassociative potentiation at a parallel fiber synapse. *Neuron* 27:611–22

Hashimotodani Y, Ohno-Shosaku T, Tsubokawa H, Ogata H, Emoto K, et al. 2005. Phospholipase Cbeta serves as a coincidence detector through its Ca^{2+} dependency for triggering retrograde endocannabinoid signal. *Neuron* 45:257–68

Hashimotodani Y, Ohno-Shosaku T, Watanabe M, Kano M. 2007. Roles of phospholipase C{beta} and NMDA receptor in activity-dependent endocannabinoid release. *J. Physiol.* 584:373–80

Hausser M, Mel B. 2003. Dendrites: bug or feature? *Curr. Opin. Neurobiol.* 13:372–83

Hebb DO. 1949. *The Organization of Behavior; A Neuropsychological Theory*. New York: Wiley. xix, 335 pp.

Hensch TK. 2005. Critical period plasticity in local cortical circuits. *Nat. Rev. Neurosci.* 6:877–88

Hirsch JC, Barrionuevo G, Crepel F. 1992. Homo- and heterosynaptic changes in efficacy are expressed in prefrontal neurons: an in vitro study in the rat. *Synapse* 12:82–85

Hoffman DA, Johnston D. 1998. Downregulation of transient K^+ channels in dendrites of hippocampal CA1 pyramidal neurons by activation of PKA and PKC. *J. Neurosci.* 18:3521–28

Hoffman DA, Johnston D. 1999. Neuromodulation of dendritic action potentials. *J. Neurophysiol.* 81:408–11

Hoffman DA, Magee JC, Colbert CM, Johnston D. 1997. K^+ channel regulation of signal propagation in dendrites of hippocampal pyramidal neurons. *Nature* 387:869–75

Holmgren CD, Zilberter Y. 2001. Coincident spiking activity induces long-term changes in inhibition of neocortical pyramidal cells. *J. Neurosci.* 21:8270–77

Humeau Y, Luthi A. 2007. Dendritic calcium spikes induce bi-directional synaptic plasticity in the lateral amygdala. *Neuropharmacology* 52:234–43

Iriki A, Pavlides C, Keller A, Asanuma H. 1989. Long-term potentiation in the motor cortex. *Science* 245:1385–87

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Ito M. 2005. Bases and implications of learning in the cerebellum—adaptive control and internal model mechanism. *Prog. Brain Res.* 148:95–109

Jacob V, Brasier DJ, Erchova I, Feldman D, Shulz DE. 2007. Spike timing-dependent synaptic depression in the in vivo barrel cortex of the rat. *J. Neurosci.* 27:1271–84

Ji D, Wilson MA. 2007. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* 10:100–7

Johnston D, Hoffman DA, Colbert CM, Magee JC. 1999. Regulation of back-propagating action potentials in hippocampal neurons. *Curr. Opin. Neurobiol.* 9:288–92

Kampa BM, Clements J, Jonas P, Stuart GJ. 2004. Kinetics of Mg^{2+} unblock of NMDA receptors: implications for spike-timing dependent synaptic plasticity. *J. Physiol.* 556:337–45

Kampa BM, Letzkus JJ, Stuart GJ. 2006. Requirement of dendritic calcium spikes for induction of spike-timing-dependent synaptic plasticity. *J. Physiol.* 574:283–90

Karmarkar UR, Buonomano DV. 2002. A model of spike-timing dependent plasticity: one or two coincidence detectors? *J. Neurophysiol.* 88:507–13

Karmarkar UR, Dan Y. 2006. Experience-dependent plasticity in adult visual cortex. *Neuron* 52:577–85

Katz LC, Shatz CJ. 1996. Synaptic activity and the construction of cortical circuits. *Science* 274:1133–38

Kilgard MP, Merzenich MM. 1998. Cortical map reorganization enabled by nucleus basalis activity. *Science* 279:1714–18

Kirkwood A, Bear MF. 1994. Hebbian synapses in visual cortex. *J. Neurosci.* 14:1634–45

Kobayashi K, Poo MM. 2004. Spike train timing-dependent associative modification of hippocampal CA3 recurrent synapses by mossy fibers. *Neuron* 41:445–54

Koester HJ, Sakmann B. 1998. Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials. *Proc. Natl. Acad. Sci. USA* 95:9596–601

Larkum ME, Kaiser KM, Sakmann B. 1999a. Calcium electrogenesis in distal apical dendrites of layer 5 pyramidal cells at a critical frequency of back-propagating action potentials. *Proc. Natl. Acad. Sci. USA* 96:14600–4

Larkum ME, Zhu JJ, Sakmann B. 1999b. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* 398:338–41

Lee I, Rao G, Knierim JJ. 2004. A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. *Neuron* 42:803–15

Lester RA, Clements JD, Westbrook GL, Jahr CE. 1990. Channel kinetics determine the time course of NMDA receptor-mediated synaptic currents. *Nature* 346:565–67

Letzkus JJ, Kampa BM, Stuart GJ. 2006. Learning rules for spike timing-dependent plasticity depend on dendritic synapse location. *J. Neurosci.* 26:10420–29

Levy WB, Steward O. 1983. Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8:791–97

Li CY, Lu JT, Wu CP, Duan SM, Poo MM. 2004. Bidirectional modification of presynaptic neuronal excitability accompanying spike timing-dependent synaptic plasticity. *Neuron* 41:257–68

Lin YW, Min MY, Chiu TH, Yang HW. 2003. Enhancement of associative long-term potentiation by activation of beta-adrenergic receptors at CA1 synapses in rat hippocampal slices. *J. Neurosci.* 23:4173–81

Linden DJ, Connor JA. 1995. Long-term synaptic depression. *Annu. Rev. Neurosci.* 18:319–57

Lisman J. 1989. A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA* 86:9574–78

Liu L, Wong TP, Pozza MF, Lingenhoehl K, Wang Y, et al. 2004. Role of NMDA receptor subtypes in governing the direction of hippocampal synaptic plasticity. *Science* 304:1021–24
Liu QS, Pu L, Poo MM. 2005. Repeated cocaine exposure in vivo facilitates LTP induction in midbrain dopamine neurons. *Nature* 437:1027–31
Louie K, Wilson MA. 2001. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29:145–56
Maejima T, Oka S, Hashimotodani Y, Ohno-Shosaku T, Aiba A, et al. 2005. Synaptically driven endocannabinoid release requires Ca^{2+}-assisted metabotropic glutamate receptor subtype 1 to phospholipase Cbeta4 signaling cascade in the cerebellum. *J. Neurosci.* 25:6826–35
Magee JC, Johnston D. 1997. A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275:209–13
Magee JC, Johnston D. 2005. Plasticity of dendritic function. *Curr. Opin. Neurobiol.* 15:334–42
Malenka RC, Bear MF. 2004. LTP and LTD: an embarrassment of riches. *Neuron* 44:5–21
Markram H, Lubke J, Frotscher M, Sakmann B. 1997. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275:213–15
Martin SJ, Grimwood PD, Morris RG. 2000. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu. Rev. Neurosci.* 23:649–711
Mayer ML, Westbrook GL, Guthrie PB. 1984. Voltage-dependent block by Mg^{2+} of NMDA responses in spinal cord neurones. *Nature* 309:261–63
Mehta MR, Barnes CA, McNaughton BL. 1997. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* 94:8918–21
Mehta MR, Quirk MC, Wilson MA. 2000. Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25:707–15
Meliza CD, Dan Y. 2006. Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking. *Neuron* 49:183–89
Meredith RM, Floyer-Lea AM, Paulsen O. 2003. Maturation of long-term potentiation induction rules in rodent hippocampus: role of GABAergic inhibition. *J. Neurosci.* 23:11142–46
Migliore M, Shepherd GM. 2002. Emerging rules for the distributions of active dendritic conductances. *Nat. Rev. Neurosci.* 3:362–70
Miles R, Toth K, Gulyas AI, Hajos N, Freund TF. 1996. Differences between somatic and dendritic inhibition in the hippocampus. *Neuron* 16:815–23
Mu Y, Poo MM. 2006. Spike timing-dependent LTP/LTD mediates visual experience-dependent plasticity in a developing retinotectal system. *Neuron* 50:115–25
Mulkey RM, Malenka RC. 1992. Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus. *Neuron* 9:967–75
Nadasdy Z, Hirase H, Czurko A, Csicsvari J, Buzsaki G. 1999. Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* 19:9497–507
Nevian T, Sakmann B. 2004. Single spine Ca^{2+} signals evoked by coincident EPSPs and back-propagating action potentials in spiny stellate cells of layer 4 in the juvenile rat somatosensory barrel cortex. *J. Neurosci.* 24:1689–99
Nevian T,. Sakmann B. 2006. Spine Ca^{2+} signaling in spike-timing-dependent plasticity. *J. Neurosci.* 26:11001–13
Nishida S, Johnston A. 1999. Influence of motion signals on the perceived position of spatial pattern. *Nature* 397:610–12
Nishiyama M, Hong K, Mikoshiba K, Poo MM, Kato K. 2000. Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature* 408:584–88
Nowak L, Bregestovski P, Ascher P, Herbet A, Prochiantz A. 1984. Magnesium gates glutamate-activated channels in mouse central neurones. *Nature* 307:462–65

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Pike FG, Meredith RM, Olding AW, Paulsen O. 1999. Rapid report: postsynaptic bursting is essential for "Hebbian" induction of associative long-term potentiation at excitatory synapses in rat hippocampus. *J. Physiol.* 518(Pt. 2):571–76

Pu L, Liu QS, Poo MM. 2006. BDNF-dependent synaptic sensitization in midbrain dopamine neurons after cocaine withdrawal. *Nat. Neurosci.* 9:605–7

Rall W. 1967. Distinguishing theoretical synaptic potentials computed for different soma-dendritic distributions of synaptic input. *J. Neurophysiol.* 30:1138–68

Ramachandran VS, Anstis SM. 1990. Illusory displacement of equiluminous kinetic edges. *Perception* 19:611–16

Rao RP, Sejnowski TJ. 2001a. Predictive learning of temporal sequences in recurrent neocortical circuits. *Novartis Found Symp.* 239:208–29; discussion 29–40

Rao RP, Sejnowski TJ. 2001b. Spike-timing-dependent Hebbian plasticity as temporal difference learning. *Neural Comput.* 13:2221–37

Ribeiro S, Gervasoni D, Soares ES, Zhou Y, Lin SC, et al. 2004. Long-lasting novelty-induced neuronal reverberation during slow-wave sleep in multiple forebrain areas. *PLoS Biol.* 2:E24

Rosenmund C, Feltz A, Westbrook GL. 1995. Calcium-dependent inactivation of synaptic NMDA receptors in hippocampal neurons. *J. Neurophysiol.* 73:427–30

Schuett S, Bonhoeffer T, Hubener M. 2001. Pairing-induced changes of orientation maps in cat visual cortex. *Neuron* 32:325–37

Seol GH, Ziburkus J, Huang S, Song L, Kim IT, et al. 2007. Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. *Neuron* 55:919–29

Shadlen MN, Newsome WT. 1994. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4:569–79

Shouval HZ, Bear MF, Cooper LN. 2002. A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA* 99:10831–36

Shu Y, Hasenstaub A, Badoual M, Bal T, McCormick DA. 2003. Barrages of synaptic activity control the gain and sensitivity of cortical neurons. *J. Neurosci.* 23:10388–401

Siegelbaum SA, Kandel ER. 1991. Learning-related synaptic plasticity: LTP and LTD. *Curr. Opin. Neurobiol.* 1:113–20

Sjostrom PJ, Hausser M. 2006. A cooperative switch determines the sign of synaptic plasticity in distal dendrites of neocortical pyramidal neurons. *Neuron* 51:227–38

Sjostrom PJ, Turrigiano GG, Nelson SB. 2001. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32:1149–64

Sjostrom PJ, Turrigiano GG, Nelson SB. 2003. Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. *Neuron* 39:641–54

Snowden RJ. 1998. Shifts in perceived position following adaptation to visual motion. *Curr. Biol.* 8:1343–45

Softky WR, Koch C. 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13:334–50

Spruston N, Schiller Y, Stuart G, Sakmann B. 1995. Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science* 268:297–300

Stanton PK, Sejnowski TJ. 1989. Associative long-term depression in the hippocampus induced by Hebbian covariance. *Nature* 339:215–18

Stefan K, Kunesch E, Benecke R, Cohen LG, Classen J. 2002. Mechanisms of enhancement of human motor cortex excitability induced by interventional paired associative stimulation. *J. Physiol.* 543:699–708

Stent GS. 1973. A physiological mechanism for Hebb's postulate of learning. *Proc. Natl. Acad. Sci. USA* 70:997–1001

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Stuart G, Schiller J, Sakmann B. 1997a. Action potential initiation and propagation in rat neocortical pyramidal neurons. *J. Physiol.* 505(Pt. 3):617–32

Stuart G, Spruston N, Sakmann B, Hausser M. 1997b. Action potential initiation and backpropagation in neurons of the mammalian CNS. *Trends Neurosci* 20:125–31

Stuart GJ, Hausser M. 2001. Dendritic coincidence detection of EPSPs and action potentials. *Nat. Neurosci.* 4:63–71

Stuart GJ, Sakmann B. 1994. Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature* 367:69–72

Tong G, Shepherd D, Jahr CE. 1995. Synaptic desensitization of NMDA receptors by calcineurin. *Science* 267:1510–12

Tsubokawa H, Ross WN. 1996. IPSPs modulate spike backpropagation and associated $[Ca^{2+}]i$ changes in the dendrites of hippocampal CA1 pyramidal neurons. *J. Neurophysiol.* 76:2896–906

Tsubokawa H, Ross WN. 1997. Muscarinic modulation of spike backpropagation in the apical dendrites of hippocampal CA1 pyramidal neurons. *J. Neurosci.* 17:5782–91

Tzounopoulos T, Kim Y, Oertel D, Trussell LO. 2004. Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nat. Neurosci.* 7:719–25

Tzounopoulos T, Rubio ME, Keen JE, Trussell LO. 2007. Coactivation of pre- and postsynaptic signaling mechanisms determines cell-specific spike-timing-dependent plasticity. *Neuron* 54:291–301

Vislay-Meltzer RL, Kampff AR, Engert F. 2006. Spatiotemporal specificity of neuronal activity directs the modification of receptive fields in the developing retinotectal system. *Neuron* 50:101–14

Wang HX, Gerkin RC, Nauen DW, Bi GQ. 2005. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.* 8:187–93

Wang Z, Xu NL, Wu CP, Duan S, Poo MM. 2003. Bidirectional changes in spatial dendritic integration accompanying long-term synaptic modifications. *Neuron* 37:463–72

Watanabe S, Hoffman DA, Migliore M, Johnston D. 2002. Dendritic K^+ channels contribute to spike-timing dependent long-term potentiation in hippocampal pyramidal neurons. *Proc. Natl. Acad. Sci. USA* 99:8366–71

Waters J, Helmchen F. 2004. Boosting of action potential backpropagation by neocortical network activity in vivo. *J. Neurosci.* 24:11127–36

Waters J, Schaefer A, Sakmann B. 2005. Backpropagating action potentials in neurones: measurement, mechanisms and potential functions. *Prog. Biophys. Mol. Biol.* 87:145–70

Whitaker D, McGraw PV, Pearson S. 1999. Non-veridical size perception of expanding and contracting objects. *Vision Res.* 39:2999–3009

Williams SR, Stuart GJ. 2000. Backpropagation of physiological spike trains in neocortical pyramidal neurons: implications for temporal coding in dendrites. *J. Neurosci.* 20:8238–46

Wilson MA, McNaughton BL. 1994. Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–79

Wittenberg GM, Wang SS. 2006. Malleability of spike-timing-dependent plasticity at the CA3-CA1 synapse. *J. Neurosci.* 26:6610–17

Wolters A, Sandbrink F, Schlottmann A, Kunesch E, Stefan K, et al. 2003. A temporally asymmetric Hebbian rule governing plasticity in the human motor cortex. *J. Neurophysiol.* 89:2339–45

Wolters A, Schmidt A, Schramm A, Zeller D, Naumann M, et al. 2005. Timing-dependent plasticity in human primary somatosensory cortex. *J. Physiol.* 565:1039–52

Woodin MA, Ganguly K, Poo MM. 2003. Coincident pre- and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in Cl- transporter activity. *Neuron* 39:807–20

Xu NL, Ye CQ, Poo MM, Zhang XH. 2006. Coincidence detection of synaptic inputs is facilitated at the distal dendrites after long-term potentiation induction. *J. Neurosci.* 26:3002–9

Yang SN, Tang YG, Zucker RS. 1999. Selective induction of LTP and LTD by postsynaptic $[Ca^{2+}]i$ elevation. *J. Neurophysiol.* 81:781–87

Yao H, Dan Y. 2001. Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron* 32:315–23

Yao H, Shen Y, Dan Y. 2004. Intracortical mechanism of stimulus-timing-dependent plasticity in visual cortical orientation tuning. *Proc. Natl. Acad. Sci. USA* 101:5081–86

Yao H, Shi L, Han F, Gao H, Dan Y. 2007. Rapid learning in cortical coding of visual scenes. *Nat. Neurosci.* 10:772–78

Yu X, Yoganarasimha D, Knierim JJ. 2006. Backward shift of head direction tuning curves of the anterior thalamus: comparison with CA1 place fields. *Neuron* 52:717–29

Zhang LI, Tao HW, Holt CE, Harris WA, Poo M. 1998. A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395:37–44

Zhang W, Linden DJ. 2003. The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.* 4:885–900

Zhou Q, Tao HW, Poo MM. 2003. Reversal and stabilization of synaptic modifications in a developing visual system. *Science* 300:1953–57

Zucker RS, Regehr WG. 2002. Short-term synaptic plasticity. *Annu. Rev. Physiol.* 64:355–405

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Annual Review of Neuroscience

Volume 31, 2008

Contents

Cerebellum-Like Structures and Their Implications for Cerebellar Function
Curtis C. Bell, Victor Han, and Nathaniel B. Sawtell 1

Spike Timing–Dependent Plasticity: A Hebbian Learning Rule
Natalia Caporale and Yang Dan 25

Balancing Structure and Function at Hippocampal Dendritic Spines
Jennifer N. Bourne and Kristen M. Harris 47

Place Cells, Grid Cells, and the Brain's Spatial Representation System
Edvard I. Moser, Emilio Kropff, and May-Britt Moser 69

Mitochondrial Disorders in the Nervous System
Salvatore DiMauro and Eric A. Schon 91

Vestibular System: The Many Facets of a Multimodal Sense
Dora E. Angelaki and Kathleen E. Cullen 125

Role of Axonal Transport in Neurodegenerative Diseases
Kurt J. De Vos, Andrew J. Grierson, Steven Ackerley, and Christopher C.J. Miller ... 151

Active and Passive Immunotherapy for Neurodegenerative Disorders
David L. Brody and David M. Holtzman 175

Descending Pathways in Motor Control
Roger N. Lemon 195

Task Set and Prefrontal Cortex
Katsuyuki Sakai 219

Multiple Sclerosis: An Immune or Neurodegenerative Disorder?
Bruce D. Trapp and Klaus-Armin Nave 247

Multifunctional Pattern-Generating Circuits
K.L. Briggman and W.B. Kristan, Jr. 271

Retinal Axon Growth at the Optic Chiasm: To Cross or Not to Cross
Timothy J. Petros, Alexandra Rebsam, and Carol A. Mason 295

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Brain Circuits for the Internal Monitoring of Movements
Marc A. Sommer and Robert H. Wurtz 317

Wnt Signaling in Neural Circuit Assembly
Patricia C. Salinas and Yimin Zou 339

Habits, Rituals, and the Evaluative Brain
Ann M. Graybiel 359

Mechanisms of Self-Motion Perception
Kenneth H. Britten 389

Mechanisms of Face Perception
Doris Y. Tsao and Margaret S. Livingstone 411

The Prion's Elusive Reason for Being
Adriano Aguzzi, Frank Baumann, and Juliane Bremer 439

Mechanisms Underlying Development of Visual Maps and Receptive Fields
Andrew D. Huberman, Marla B. Feller, and Barbara Chapman 479

Neural Substrates of Language Acquisition
Patricia K. Kuhl and Maritza Rivera-Gaxiola 511

Axon-Glial Signaling and the Glial Support of Axon Function
Klaus-Armin Nave and Bruce D. Trapp 535

Signaling Mechanisms Linking Neuronal Activity to Gene Expression and Plasticity of the Nervous System
Steven W. Flavell and Michael E. Greenberg 563

Indexes

Cumulative Index of Contributing Authors, Volumes 22–31 591

Cumulative Index of Chapter Titles, Volumes 22–31 595

Errata

An online log of corrections to *Annual Review of Neuroscience* articles may be found at http://neuro.annualreviews.org/

Annu. Rev. Neurosci. 2008.31:25-46. Downloaded from arjournals.annualreviews.org
by University of California - Berkeley on 10/23/09. For personal use only.

Yang Dan and Mu-Ming Poo
Physiol Rev 86:1033-1048, 2006. doi:10.1152/physrev.00030.2005

You might find this additional information useful...

This article cites 118 articles, 45 of which you can access free at:
http://physrev.physiology.org/cgi/content/full/86/3/1033#BIBL

This article has been cited by 33 other HighWire hosted articles, the first 5 are:

Human motor associative plasticity induced by paired bihemispheric stimulation
S. Koganemaru, T. Mima, M. Nakatsuka, Y. Ueki, H. Fukuyama and K. Domen
J. Physiol., October 1, 2009; 587 (19): 4629-4644.
[Abstract] [Full Text] [PDF]

A Selective Interplay between Aberrant EPSPKA and INaP Reduces Spike Timing Precision in Dentate Granule Cells of Epileptic Rats
J. Epsztein, E. Sola, A. Represa, Y. Ben-Ari and V. Crepel
Cereb Cortex, September 9, 2009; 0 (2009): bhp156v2-bhp156.
[Abstract] [Full Text] [PDF]

Inhibition and Recurrent Excitation in a Computational Model of Sparse Bursting in Song Nucleus HVC
L. Gibb, T. Q. Gentner and H. D. I. Abarbanel
J Neurophysiol, September 1, 2009; 102 (3): 1748-1762.
[Abstract] [Full Text] [PDF]

Memory Retention and Spike-Timing-Dependent Plasticity
G. Billings and M. C. W. van Rossum
J Neurophysiol, June 1, 2009; 101 (6): 2775-2788.
[Abstract] [Full Text] [PDF]

Spike-Timing-Dependent Synaptic Plasticity and Synaptic Democracy in Dendrites
A. Gidon and I. Segev
J Neurophysiol, June 1, 2009; 101 (6): 3226-3234.
[Abstract] [Full Text] [PDF]

Medline items on this article's topics can be found at http://highwire.stanford.edu/lists/artbytopic.dtl on the following topics:
Physiology .. Dendrites
Physiology .. Neuronal Excitability
Veterinary Science .. Long-Term Depression
Physiology .. Long-Term Potentiation
Physiology .. Neural Circuit
Physiology .. Humans

Updated information and services including high-resolution figures, can be found at:
http://physrev.physiology.org/cgi/content/full/86/3/1033

Additional material and information about ***Physiological Reviews*** can be found at:
http://www.the-aps.org/publications/prv

This information is current as of October 23, 2009 .

Physiological Reviews provides state of the art coverage of timely issues in the physiological and biomedical sciences. It is published quarterly in January, April, July, and October by the American Physiological Society, 9650 Rockville Pike, Bethesda MD 20814-3991. Copyright © 2005 by the American Physiological Society. ISSN: 0031-9333, ESSN: 1522-1210. Visit our website at http://www.the-aps.org/.

Downloaded from physrev.physiology.org on October 23, 2009

Physiol Rev 86: 1033–1048, 2006;
doi:10.1152/physrev.00030.2005.

Spike Timing-Dependent Plasticity: From Synapse to Perception

YANG DAN AND MU-MING POO

Division of Neurobiology, Department of Molecular and Cell Biology, and Helen Wills Neuroscience Institute, University of California, Berkeley, California

I. Introduction 1033
II. Cellular Mechanisms Underlying Synaptic Spike Timing-Dependent Plasticity 1034
A. LTP window 1035
B. LTD window 1035
C. Modulation of STDP by inhibitory inputs 1036
III. Spike Timing-Dependent Plasticity With Complex Spatiotemporal Activity Patterns 1036
A. Dependence on dendritic location 1036
B. Complex spike trains 1037
IV. Spike Timing-Dependent Plasticity in Vivo 1038
A. Electrical stimulation 1038
B. Sensory stimulation 1039
C. Natural stimuli 1039
D. Persistence of STDP in vivo 1041
V. Nonsynaptic Aspects of Spike Timing-Dependent Plasticity 1042
A. STDP of intrinsic neuronal excitability 1042
B. STDP of local dendritic excitability and synaptic integration 1043
VI. Spread of Synaptic Spike Timing-Dependent Plasticity in Neural Circuits 1043
A. Heterosynaptic effects of LTP/LTD induction 1043
B. Spread of LTP/LTD in cell culture 1044
C. Spread of LTP/LTD in vivo 1045
VII. Concluding Remarks 1045

Dan, Yang, and Mu-Ming Poo. Spike Timing-Dependent Plasticity: From Synapse to Perception. *Physiol Rev* 86: 1033–1048, 2006; doi:10.1152/physrev.00030.2005.—Information in the nervous system may be carried by both the rate and timing of neuronal spikes. Recent findings of spike timing-dependent plasticity (STDP) have fueled the interest in the potential roles of spike timing in processing and storage of information in neural circuits. Induction of long-term potentiation (LTP) and long-term depression (LTD) in a variety of in vitro and in vivo systems has been shown to depend on the temporal order of pre- and postsynaptic spiking. Spike timing-dependent modification of neuronal excitability and dendritic integration was also observed. Such STDP at the synaptic and cellular level is likely to play important roles in activity-induced functional changes in neuronal receptive fields and human perception.

I. INTRODUCTION

Since the discovery of persistent enhancement of synaptic transmission by tetanic stimulation in the hippocampus (14), a phenomenon now generally referred to as long-term potentiation (LTP), the study of activity-dependent synaptic plasticity has become one of the most active areas in neurobiology (66, 68). Two features of LTP, the associativity and input specificity, match the properties of some forms of learning and memory, suggesting that LTP may underlie such cognitive functions. Traditionally, LTP is induced by high-frequency presynaptic stimulation or by pairing low-frequency stimulation with postsynaptic depolarization. Prolonged low-frequency stimulation was also found to induce long-term depression (LTD) (53, 77). Thus synaptic efficacy can be modified in a bidirectional manner.

In studying the temporal specificity of associative synaptic modification in the hippocampus, a region known to be important for memory formation, Levy and

0031-9333/06 $18.00 Copyright © 2006 the American Physiological Society

Downloaded from physrev.physiology.org on October 23, 2009

Steward (57) noted that when a weak and a strong input from entorhinal cortex to the dentate gyrus were activated together, the temporal order of activation was crucial. LTP of the weak input was induced when the strong input was activated concurrently with the weak input or following it by as much as 20 ms. Interestingly, LTD was induced when the temporal order was reversed. Later studies have further addressed the importance of the temporal order of pre- and postsynaptic spiking in long-term modification of a variety of glutamatergic synapses and have defined the "critical windows" for spike timing (12, 13, 15, 24, 28, 31, 37, 65, 67, 97, 107, 119). As illustrated in Figure 1, when presynaptic spiking precedes postsynaptic spiking (hereafter referred to as "pre-post") within a window of several tens of milliseconds, LTP is induced, whereas spiking of the reverse order ("post-pre") leads to LTD. This form of activity-dependent LTP/LTD is now referred to as spike timing-dependent plasticity (STDP) (99). While the dependence of synaptic modification on the pre/post spike order is commonly observed, the width of the STDP window varies. In Table 1, we list the pre/post spike intervals at which LTP and LTD have been detected at different synapses. In addition to these glutamatergic synapses, a recent study of GABAergic synapses in rat hippocampal cultures and slices showed that repetitive postsynaptic spiking within 20 ms either before or after presynaptic activation led to a persistent change in the GABAergic synaptic strength (115), giving rise to a symmetric temporal window.

Beyond the initial characterization of the STDP windows, more recent experimental studies have addressed the following questions: What cellular mechanisms underlie the temporal windows for LTP/LTD? What are the rules determining the synaptic modification induced by complex spike patterns, including high-frequency bursts? How does STDP affect the operation of neural circuits and higher brain functions such as sensory perception? In addition to synaptic modification, is there STDP in neuronal excitability and dendritic integration? We here review findings pertinent to these questions.

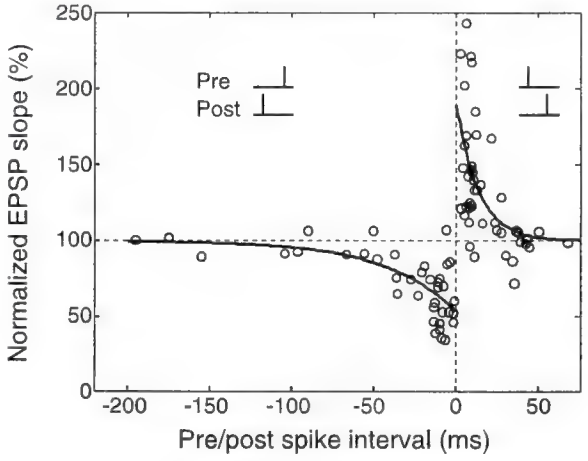

FIG. 1. Synaptic modification induced by repetitively paired pre- and postsynaptic spikes in layer 2/3 of visual cortical slices from the rat. Each symbol represents result from one experiment. Curves are single exponential, least-squares fits of the data. *Insets* depict the sequence of spiking in the pre- and postsynaptic neurons. EPSP, excitatory postsynaptic potential.

II. CELLULAR MECHANISMS UNDERLYING SYNAPTIC SPIKE TIMING-DEPENDENT PLASTICITY

In conventional protocols using steady postsynaptic depolarization, high-frequency presynaptic stimulation induces LTP and low-frequency stimulation induces LTD, but in STDP low-frequency stimulation can be used to induce both LTP and LTD. While in both types of protocols activation of *N*-methyl-D-aspartate (NMDA) subtype of glutamate receptors (NMDARs) (e.g., Refs. 13, 24, 65, 67, 119) and elevation of postsynaptic Ca^{2+} level are required, the effectiveness of postsynaptic spiking and steady depolarization in achieving the required Ca^{2+} level is likely to be different. This may account for the higher efficiency of the spike-timing protocol in long-term modification of excitatory synapses in hippocampal cultures (13) and midbrain slices (60). The NMDARs are largely blocked by Mg^{2+} at hyperpolarized membrane potentials, but the block can be relieved by depolarization (69, 83), leading to the idea that the NMDAR serves as the coincidence detector for pre/post activity. For LTP and LTD induced by conventional protocols, different cascades of signaling events are set in motion by high- and low-level postsynaptic Ca^{2+} elevation, respectively (66). Does the conventional Ca^{2+}-based model of LTP/LTD also explain the temporally asymmetric STDP? One would expect that pre-post spiking leads to a brief high-level Ca^{2+} influx, due to effective activation of NMDARs by postsynaptic spiking, while post-pre spiking leads to a low-level Ca^{2+} rise due to the limited extent of NMDAR activation by the afterdepolarization associated with the postsynaptic action potential (AP). Fluorescence Ca^{2+} imaging studies indeed demonstrated that Ca^{2+} influx through NMDARs and voltage-dependent Ca^{2+} channels (VDCCs) exhibits supralinear summation with pre-post spiking and sublinear summation with post-pre spiking (56, 79). In support of this Ca^{2+} model for STDP, pre-post spiking under partial inactivation of NMDARs in CA1 hippocampus leads to the induction of LTD instead of LTP (81). However, such a simple model is unlikely to provide a complete explanation for the STDP window.

Downloaded from physrev.physiology.org on October 23, 2009

TABLE 1. *Temporal window for STDP of glutamatergic synapses*

Synapse	LTP Window, ms	LTD Window, ms	Reference Nos.
Rat (hippocampal slice) entorhinal cortex → dentate gyrus	0–20	<0	57
Electric fish (electrosensory lobe, slice) parallel fiber → Purkinje-like cell	<0 or >50	0–50	12
Xenopus tadpole (in vivo) retina → tectum	0–40	0 to −20	119
Zebra finch (brain slice) $LMAN_R$ → LMAN	0–15	0 to −7	15
Rat (hippocampal culture)	0–30	0 to −40	13, 58
Rat (hippocampal slice culture) CA3 → CA3	15	0 to −200	24
Rat (hippocampal slice) Schaffer collateral/commissural → CA1	0–10	−15 to −30 and 15–25	81
Rat (hippocampal slice) Schaffer collateral → CA1	3–10*		59
Rat (neocortical slice) L5 → L5	10	−10	67
Rat (barrel cortical slice) L4 → L4	NF	−10 to 25	28
Rat (barrel cortical slice) L4 → L2/3	0–15	0 to −100	31
Rat (visual cortical slice) L5 → L5	10	0 to −50	97
Rat (visual cortical slice) input to L2/3; proximal (A), distal (B)	0–25 (A) 0–25 (B)	0 to −50 (A) 0 to −100 (B)	38
Mouse (brain stem slice) parallel fiber → fusiform (A), cartwheel (B)	5–10 (A) NF (B)	−5 (A) 5 (B)	107

Due to the sparsity of data points in some studies, the listed range of intervals may not reflect the full width of the window. For window data: positive interval, pre-post; negative interval, post-pre. NF, not found. * Activation of β-adrenergic receptors broadens this window to 3–15 ms.

A. LTP Window

The ~20 ms pre-post window for LTP induction is much shorter than the time constant for the dissociation of glutamate from NMDARs (109). One explanation of this narrow LTP window is the kinetics of Mg^{2+} unblock of the NMDARs. Using nucleated patches of neocortical pyramidal neurons, Kampa et al. (50) measured the rate of depolarization-induced Mg^{2+} unblock of NMDARs at different times after a brief pulse of glutamate application. They found that Mg^{2+} unblock consists of a fast and a slow component, whose relative amplitudes depend on the timing of the depolarization relative to the glutamate pulse, with the fast component preferentially reduced at later times. Thus the postsynaptic spikes arriving immediately after glutamate binding are more effective in opening NMDARs, thus sharpening the time window for LTP induction. In addition to the property of NMDARs, postsynaptic cytoplasmic processes, including Ca^{2+}-dependent Ca^{2+} release from internal stores and activation of downstream effectors, both of which are likely to be highly nonlinear, may further boost the effect of the earlier arriving postsynaptic APs and sharpen the LTP window. The importance of cytoplasmic processes is further supported by the finding that activation of β-adrenergic receptors can increase the width (but not the magnitude) of the LTP window at Schaffer collateral-CA1 pyramidal cell synapses (59), an effect that depends on protein kinase A or mitogen-activated protein kinase. Neuromodulatory inputs may also affect the STDP window by modifying the AP profile through the modulation of the density and properties of axonal and dendritic ion channels, such as transient (I_A) and Ca^{2+}-activated (BK and SK) K^+ channels (95, 112).

B. LTD Window

The simplest explanation for the post-pre LTD window is that the afterdepolarization associated with the AP causes a partial opening of NMDARs, resulting in a low-level Ca^{2+} influx required for LTD induction. However, such a simple model inevitably predicts the existence of a LTD window at positive (pre-post) intervals longer than those for LTP induction. This is because, as the amount of Ca^{2+} influx through NMDARs decreases with the pre-post interval, it must pass through the range appropriate for LTD induction before reaching the baseline. Although such a pre-post LTD window has been found in hippocampal slices (81), it was not reported in any other study. A potential resolution of this paradox is that excitatory postsynaptic potentials (EPSPs) in the absence of postsynaptic AP may already induce Ca^{2+} elevation above that required for LTD induction (but insufficient to induce LTP); thus pre-post spiking can only lead to LTP. For post-pre spiking, other mechanisms such as spiking-induced afterhyperpolarization or NMDAR desensitization (88, 106, 108) reduce the Ca^{2+} influx through NMDARs to a level appropriate for LTD induction (38). Alternatively, a second coincidence detector independent of NMDAR is responsible for LTD at post-pre intervals. Karmarkar and Buonomano (51) suggest that the mGluR pathway, which is coupled to Ca^{2+} influx through VDCCs, is a likely candidate for this second coincidence detector. This model is consistent with the finding of a mGluR-dependent form of homosynaptic LTD in the hippocampus induced by both low-frequency stimulation (47, 84) and post-pre spiking (82). Furthermore, postsynaptic spiking-triggered secretion of retrograde factors, e.g., endocannabinoids, may act together with the activation of presyn-

Downloaded from physrev.physiology.org on October 23, 2009

aptic NMDARs to induce LTD (96), as blocking the degradation of endocannabinoids was found to widen the LTD window. Besides Ca^{2+} influx through NMDARs, activation of L-type VDCCs by postsynaptic spiking (13) and Ca^{2+} release from internal stores (81) are also shown to be necessary for spike-timing dependent LTD. Because Ca^{2+} from these sources and through NMDARs may have different spatiotemporal profiles, they may play different roles in the induction of synaptic modifications.

C. Modulation of STDP by Inhibitory Inputs

Because both feed-forward and feedback GABAergic inhibition will affect dendritic depolarization induced by excitatory inputs and thus modulate postsynaptic Ca^{2+} elevation, the presence of correlated GABAergic inputs onto the postsynaptic neuron may profoundly influence the induction of LTP/LTD (113). During development, the Schaffer collateral input to CA1 pyramidal neurons becomes progressively less susceptible to LTP induction by pre/post spike pairing (74), and postsynaptic spike bursts are necessary in adult animals to induce LTP. This developmental reduction in the effectiveness of spike pairing for LTP induction can be reversed by blocking $GABA_A$ receptor-mediated inhibition. Similar GABA-mediated "gating" of LTP is also found at excitatory inputs to dopamine neurons in the ventral tegmental area (VTA) of the rat midbrain (60). Interestingly, repeated exposure of the rat to cocaine in vivo leads to a reduction of GABAergic inhibition, allowing LTP induction by the spike pairing protocol in acute midbrain slices, while enhancing GABAergic transmission in the slices from cocaine-treated animal with diazepam eliminated LTP induction. By titrating the strength of GABAergic inputs with different doses of a $GABA_A$ antagonist, Liu et al. (60) concluded that a 30% reduction of GABAergic transmission in normal rat VTA can open the gate for LTP induction. Finally, there is evidence that the LTD window at positive timing found in hippocampal slices (81) may be due to feedforward inhibition in the circuit (104).

III. SPIKE TIMING-DEPENDENT PLASTICITY WITH COMPLEX SPATIOTEMPORAL ACTIVITY PATTERNS

A. Dependence on Dendritic Location

A prominent morphological feature of the neuron is its extensive dendritic tree, where most of the synapses are located. Due to cable filtering and the nonuniform distribution of voltage-dependent ion channels (49), neuronal processing of each synaptic input depends strongly on its dendritic location (44, 94). Because STDP involves the interaction between the synaptic input and the back-propagating AP as well as the downstream cellular mechanisms highly localized at the synaptic site, it is likely to be dendritic location dependent. Indeed, both the amplitude and width of the STDP window were found to vary along the apical dendrite of the pyramidal neuron in layer 2/3 (L2/3) of visual cortical slices (38). At the intermediate-distal portion of the dendrite (100–150 μm from soma), the magnitude of LTP is smaller and the temporal window for LTD is broader than at the proximal dendrite. The smaller magnitude of LTP may be related to the attenuation of the back-propagating AP at the distal dendrite. The LTD window, on the other hand, was found to correlate with the window for the suppression of NMDA receptor-mediated EPSPs by the back-propagating APs, both across dendritic locations and under various pharmacological manipulations. This is consistent with the notion that the AP-induced suppression of EPSPs plays an important role in LTD induction, and the variability in the LTD window among different synapses may be attributed in part to differences in the dendritic location of the synapse.

Functionally, the dendritic inhomogeneity of the STDP window may allow differential selection of synaptic inputs at different portions of the dendrites according to the temporal characteristics of the presynaptic spike trains, as suggested by simulation studies (38). Although spatial inhomogeneity in the electrical properties and nonlinearity in dendritic integration are thought to endow individual pyramidal neurons with enhanced processing capacity (44, 94), the computational power can only be harvested if the inputs carrying different signals are segregated into distinct regions (8, 72). In fact, the importance of domain-specific inputs has been well recognized for inhibitory interneurons (35, 54, 86). The function and mechanism of the dendritic segregation of excitatory inputs remain largely unexplored.

While differences in the amplitude and width of the STDP window have been observed from the proximal to the intermediate segment of the apical dendrite, other differences may exist in the more distal region. For example, in some neurons, the back-propagating APs may not reach the distal tip of the dendrite, which will preclude the EPSP-AP interaction required for STDP induction. Instead, LTP may be induced by other mechanisms such as the local Ca^{2+} spikes resulting from strong synaptic inputs (43). Furthermore, while studied much less extensively than the main trunk of the apical dendrite due to technical difficulties, the basal and oblique dendrites of pyramidal neurons in fact receive most of the synaptic inputs (85). How STDP and other forms of synaptic modification operate on these dendrites remain to be investigated.

Downloaded from physrev.physiology.org on October 23, 2009

B. Complex Spike Trains

In most in vitro studies of STDP, the induction protocol consists of relatively simple spike patterns with pre/post spikes paired at regular intervals. The advantage of this approach is that each induction pattern can be described by a small number of parameters, and the dependence of synaptic modification on these parameters (e.g., pre/post spike interval) can be easily determined. For neural circuits in vivo, however, spiking in both pre- and postsynaptic cells is likely to be irregular (98), with occasional high-frequency bursts. How well does the STDP window (Fig. 1) account for the effects of complex spike trains? Do all the pre/post spike pairs contribute independently to long-term synaptic modification?

This question was addressed in L2/3 of visual cortical slices by progressively increasing the complexity of the induction pattern (37). First, a single spike, either pre- or postsynaptic, was added to the spike pair to form a "triplet," which consists of two pre/post pairs. By measuring synaptic modification at various intervals, the first spike pair was found to play a dominant role in determining the sign and magnitude of synaptic modification. Based on this observation, a simple phenomenological model was proposed in which the efficacy of each spike in synaptic modification is suppressed by the preceding spike in the same neuron in a time-dependent manner, and the contribution of each spike pair in synaptic modification is scaled by the efficacies of both spikes. Compared with the default "independent model," the suppression model significantly improved the prediction of synaptic modification induced by more complex spike patterns, including spike quadruplets and spike train segments recorded in vivo in response to natural stimuli. Interestingly, such suppressive interaction between neighboring spikes is consistent with the prediction of an adaptive STDP learning rule designed to stabilize the output firing rate of the postsynaptic neuron with changing input rate (52).

The effects of complex spike trains in LTP/LTD induction were also examined in L5 of visual cortical slices, using paired bursts of spikes with varying pre/post intervals and burst frequencies (97). In addition to the pre/post interval, synaptic modification was also found to depend on the firing frequency within each burst in a manner that cannot be accounted for by the STDP window. The magnitude of LTP, but not LTD, increases with the burst frequency (67, 97), and at high frequencies LTP was observed regardless of the pre/post interval (97). This finding suggests that LTP induced by pre-post pairs "wins over" LTD induced by post-pre pairs, and a model implementing this "LTP dominant" rule indeed improved the prediction of the effects of Poisson spike trains in synaptic modification. Alternatively, since with high-frequency bursts synaptic modification is no longer sensitive to the timing of individual spikes, perhaps a burst of spikes should be considered collectively as a basic unit in synaptic modification. Consistent with this idea, a recent study in the CA3 region of the hippocampus (55) showed that LTP of the associational/commissural connections can be induced by pairing spike bursts in the mossy fibers and the association/commissural pathway, and this effect depends on the order and interval between the pre/post bursts rather than between individual spikes.

For a synaptic learning rule used to understand computation at the circuit level, the disadvantage of treating each burst as a distinct unit is that separate rules must be characterized for these bursts. It is unclear how many types of bursts one must consider, with varying numbers of spikes and burst frequencies, and whether each type warrants a different rule. In an attempt to "rescue" the learning rule based on the timing of individual spikes, Froemke et al. (39) systematically varied the frequency, timing, and number of spikes in both the pre- and postsynaptic bursts to identify the situations under which the suppression model (37) breaks down. Consistent with the finding of Sjostrom et al. (97) in L5 synapses, L2/3 synapses also exhibit a frequency-dependent transition from LTD to LTP induced by a postsynaptic burst (5 spikes) preceding a presynaptic burst. However, this frequency dependence is a natural consequence of timing-dependent interactions among individual spikes in the paired bursts, and a simple modification of the original suppression model (37) can account for synaptic modifications induced by a variety of spike trains. Furthermore, pharmacological experiments suggest that short-term depression of presynaptic transmitter release and the kinetics of postsynaptic APs during a burst play important roles in the suppressive interactions among multiple spikes in the induction of long-term synaptic modification.

The suppression rule described above for cortical slices, however, does not seem to apply to cultured hippocampal synapses. Wang et al. (110) found that pre-post-pre triplets induce little synaptic modification, whereas post-pre-post triplets induce significant LTP. Thus the effect of the first pair appears to be "vetoed" by the second pair, opposite to that described by the suppression model (37). Given the similarity in the induction protocol, the discrepancy between the findings in cortical slice and hippocampal culture may be attributed to the different properties of synapses in these circuits, such as short-term synaptic plasticity, kinetics of postsynaptic APs, or downstream intracellular mechanisms leading to LTP and LTD. Pharmacological experiments by Wang et al. (110) suggest that the Ca^{2+}/calmodulin kinase II (CaMKII)-mediated potentiation process and a calcineurin-mediated depression process are both activated by the triplets, and they interact in a nonlinear manner.

Downloaded from physrev.physiology.org on October 23, 2009

Finally, in addition to interspike interactions at the time scale of tens of milliseconds, there are also interactions at longer time scales in STDP induction. For example, the amount and persistence of synaptic modification may depend on the temporal distribution of the spiking activity. While synaptic modification may saturate within a single induction episode, a resting period of minutes to hours may allow the synapse to recover its sensitivity to additional episodes. Such dependence on the temporal pattern of induction has been demonstrated in the developing *Xenopus* visual system (122), a phenomenon reminiscent of the differential effects of "mass" versus "spaced" learning (see below). The susceptibility of the synapse to modification may also depend on the state of the synapse, which can be a function of both the current synaptic strength (13, 23, 110) and the prior history of modification (75). Further clarification of these issues is important for understanding the functional roles of STDP in vivo.

In summary, a universally applicable model for predicting synaptic modification induced by complex spike trains must incorporate all the cellular events involved in the induction of LTP and LTD. Of particular importance is the spatiotemporal pattern of Ca^{2+} elevation that is regulated by a variety of factors. On the other hand, elucidation of simple and analytically tractable phenomenological rules is useful for exploring the functional implications of synaptic plasticity in large neuronal circuits, even in the absence of a complete understanding of the cellular mechanisms.

IV. SPIKE TIMING-DEPENDENT PLASTICITY IN VIVO

While most of the earlier studies focused on characterizing STDP at identified synapses in slices or cell cultures, more recent studies have begun to address the functional consequences of STDP in vivo. These studies are briefly summarized in Table 2. For more detailed descriptions of these studies, we have divided them into three groups, based on how spike timing is controlled experimentally to induce circuit modifications. First, postsynaptic spiking is evoked by electrical stimulation, while presynaptic activation is induced by either sensory or electrical stimulation. Second, neuronal spike timing is manipulated by pure sensory stimuli. Third, natural patterns of sensory stimuli are applied, but the resulting changes in the circuit functions are thought to involve STDP.

A. Electrical Stimulation

The first in vivo study of STDP was carried out in the developing *Xenopus* retinotectal circuit. Using whole cell recordings from the tectal neurons to monitor the synaptic inputs from retinal ganglion cells, Zhang et al. (119) found that these synapses undergo STDP with temporally asymmetric time windows, similar to those found in hippocampal cultures (13). The demonstration of STDP in this circuit in vivo has important implications in the developmental refinement of the retinotectal projection: a strong input that can elicit spiking by itself should have a competitive advantage among converging inputs, because the onset of its EPSPs relative to the postsynaptic spiking is always within the potentiation window. For a weak input not sufficient to evoke postsynaptic spiking on its own, however, the timing of its activation relative to other inputs becomes critical in determining its fate: activation within a 20-ms window before postsynaptic spiking initiated by other inputs leads to potentiation, whereas activation shortly after spiking leads to depression and perhaps the eventual elimination of the input.

The functional consequence of STDP has also been examined in the mammalian visual system. In the kitten visual cortex (92), repetitive pairing of a brief oriented visual stimulus with electrical stimulation of the cortex for 2–4 h induced substantial, long-lasting reorganization of the orientation map, as revealed by intrinsic optical imaging. At the site of stimulation, the cortical response to the paired orientation was enhanced or depressed, when the cortex was activated visually before or after the electrical stimulation by 10–20 ms, respectively, consis-

TABLE 2. *STDP of circuit functions in vivo*

Species and Circuits	Stimuli (Pre/Post)	Modification	Reference Nos.
Xenopus retinotectal system	Electrical stimulation/intracellular current injection	Synaptic response	119
Rat visual cortex	Visual stimulation/intracellular current injection	Spatiotemporal receptive field	73
Rat barrel cortex	Whisker deprivation	Synaptic response	5, 18
Rat hippocampus	Exploratory movement	Place field position	71
Kitten visual cortex	Visual stimulation/intracortical stimulation	Orientation tuning	92
Cat visual cortex	Visual stimulation/visual stimulation	Receptive field position	40, 117, 118
Cat visual cortex	Motion stimuli	Receptive field position	41
Human motor cortex	Somatosensory stimulation/TMS stimulation	TMS response	101, 114

Downloaded from physrev.physiology.org on October 23, 2009

tent with STDP of intracortical connections. In parallel with optical imaging, single-unit recordings showed that pairing at the positive interval (visual before electrical) caused an ~20° shift in the preferred orientation of the neurons toward the paired orientation. In the developing rat visual cortex, whole cell recordings showed that synaptic responses evoked by a flashed bar can be modified by pairing the bar stimulus with postsynaptic spiking elicited by the whole cell electrode (73). The direction of the modification depends on the temporal order of the synaptic response and the postsynaptic spiking in a manner consistent with STDP found in visual cortical slices (37). Note that, while earlier studies using extracellular K^+ or pharmacological agents (33, 34) to stimulate the cortical neurons demonstrated the role of coincidence between sensory stimuli and cortical activation in receptive field (RF) modification, the two recent studies described above provided more precise control of cortical spike timing at the millisecond level, thus revealing the importance of the order of visual and cortical activation in RF modification.

Remarkably, plasticity resembling STDP has been demonstrated in the human motor cortex (101, 114). The induction protocol, termed paired associative stimulation (PAS), consists of repetitive pairing of median nerve stimulation (which activates somatosensory inputs to the primary motor cortex) with transcranial magnetic stimulation (TMS) of the motor cortex. Response to the TMS stimulation, measured by electromyographic (EMG) activity of the corresponding muscle, was found to be potentiated or depressed by 90 pairs of PAS, depending on whether the somatosensory input arrived before or after motor cortical activation by TMS, consistent with the time window of STDP. The potentiation is blocked by the NMDAR antagonist dextromethorphan, and the depression is blocked by both dextromethorphan and nimodipine, a blocker of L-type VDCCs, identical to the pharmacological properties of STDP in hippocampal and cortical synapses (see above).

B. Sensory Stimulation

In several other studies, spiking of cortical neurons was evoked by visual stimulation in the absence of electrical stimulation, allowing a more physiological condition for studying the functional significance of STDP. In adult cat V1, pairing of visual stimuli at two orientations for several minutes induced a shift in the orientation tuning of cortical neurons, with the direction of the shift depending on the temporal order of the stimulus pair (117). The induction of significant shift required that the interval between the pair fall within ±40 ms, reminiscent of the temporal window for visual cortical STDP (37). This RF modification is likely to occur at the cortical level rather than in the earlier visual circuits, such as the retina or the lateral geniculate nucleus, because the shifts induced by monocular conditioning exhibit complete interocular transfer, indicating that the underlying circuit changes occur after signals from the two eyes converge. In addition, a model circuit with STDP implemented at the intracortical connections can account for both the temporal and orientation specificity of the effect (118).

With the assumption that the responses of these orientation-selective cortical neurons underlie the perception of stimulus orientation, a change in the neuronal tuning in one direction should lead to a shift in perceived orientation in the opposite direction. This prediction was tested in human psychophysical experiments using the conditioning stimuli similar to those in the electrophysiology experiments and a two alternative forced-choice task to measure perceived orientation before and after conditioning. Significant perceptual shifts were induced when the interval between the pair of stimuli was within ±40 ms, similar to the window measured for cat cortical neurons (117). Furthermore, the perceptual shift induced by monocular conditioning showed complete interocular transfer, indicating that the change is cortical in origin (118).

In addition to the orientation domain, stimulus-timing-dependent cortical modification has also been demonstrated in the space domain (40). Asynchronous visual stimuli flashed in two adjacent retinal regions were found to control the relative spike timing of two groups of cortical neurons with a precision of tens of milliseconds, and they induced modifications of intracortical connections (revealed by cross-correlation analysis) and shifts of the RF center (Fig. 2). These changes depended on the temporal order and interval between the pair of stimuli, consistent with STDP of intracortical connections. Parallel to the modifications in cat V1, asynchronous conditioning also induced shifts in perceived object position by human subjects. The similarity in the temporal window of stimulus-induced cortical modifications at multiple levels, from synapse to perception, strongly suggests that STDP and the asynchronous stimulus-induced functional changes in the visual cortex are causally related.

C. Natural Stimuli

To gain further insights into the functional implications of STDP, it is important to examine the effects of natural sensory stimuli. Of particular interest are the moving stimuli commonly found in natural scenes, since they can activate neighboring V1 neurons at short temporal intervals, thus modifying their synaptic connections through STDP. The interaction between motion stimuli

Downloaded from physrev.physiology.org on October 23, 2009

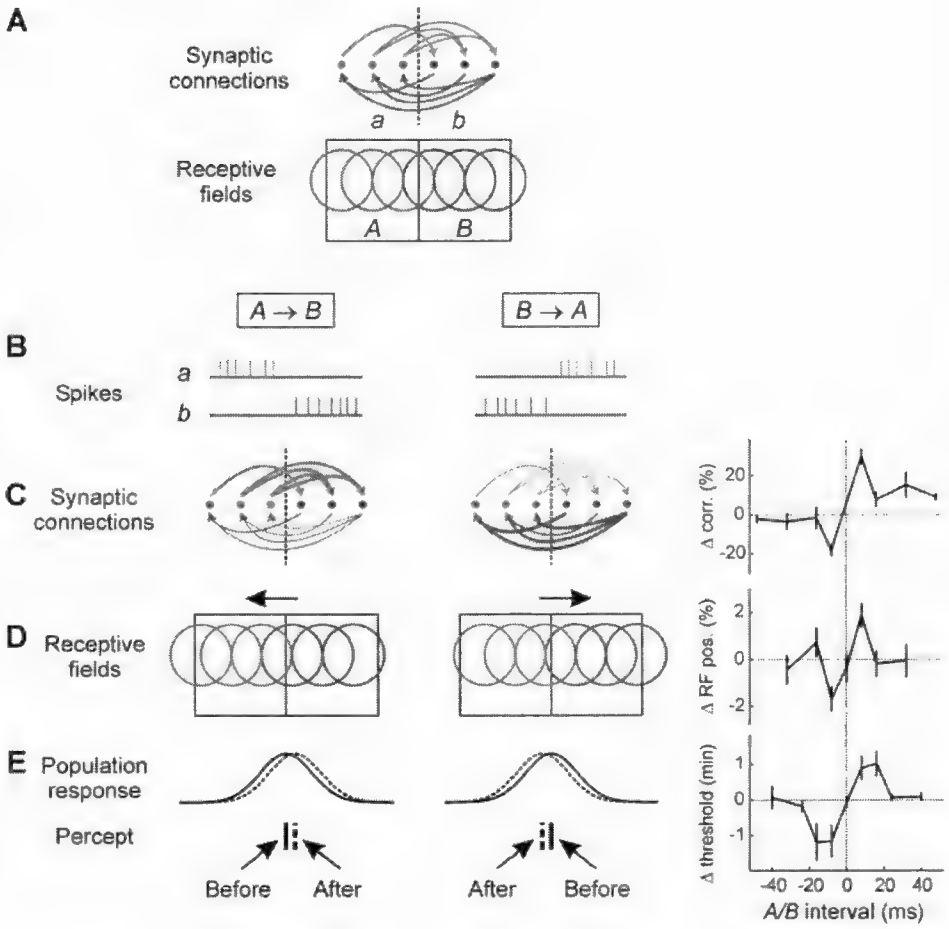

FIG. 2. A model for stimulus timing-dependent visual cortical modification. *A*: intracortical excitatory connections between two groups of neurons, *a* and *b*, whose receptive fields (RFs) fall in regions *A* and *B*, respectively. *B*: schematic illustration of neuronal spike timing in response to *A*→*B* and *B*→*A* stimulation. *C*: connection weight (represented by line thickness) after visual conditioning. *Right plot*, change in intracortical connection (from cross-correlation analysis) versus *A*/*B* interval measured in cat V1. *D*: RFs after conditioning. Arrow indicates predicted direction of shift. *Right plot*, RF shift measured in cat V1. *E*, *top*: spatial profiles of population responses evoked by a stimulus at *A*/*B* border before (solid curve) and after (dashed curve) conditioning. *Bottom*: perceived stimulus positions predicted by the model. *Right plot*, average perceptual shift measured in four human subjects. [Adapted from Fu et al. (40).]

and STDP in shaping the visual cortical function has been explored in a simple model circuit, which consists of a postsynaptic neuron ("target neuron") and two sets of direction-selective presynaptic neurons preferring either left- or rightward motion (Fig. 3). A rightward moving object should evoke a wave of spiking preferentially from the neurons preferring rightward motion ("→ neurons", red circles in Fig. 3), with the neurons on the left side of the retinotopic map firing before the target neuron, and those on the right firing afterwards. According to STDP,

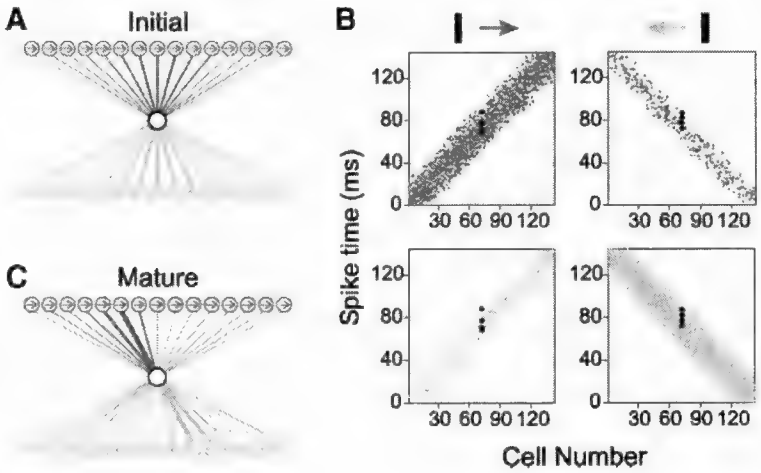

FIG. 3. Development of asymmetric visual cortical circuit through spike timing-dependent plasticity (STDP). *A*: initial circuit with symmetric connections. Arrow in each circle represents direction selectivity of the cell. Line thickness represents synaptic strength (the weakest connections are shown in gray). *B*: simulated spatiotemporal spike patterns in the circuit evoked by a small object (*top*) moving rightward (*left column*) or leftward (*right column*) across the visual field. Each red or blue dot represents a spike of a presynaptic neuron (indexed 1 to 143, from *left* to *right*) with the corresponding color. Black dots represent spikes of the target neuron. *C*: mature circuit after extensive exposure to motion stimuli in both directions, during which the connections underwent spike timing-dependent modification.

Downloaded from physrev.physiology.org on October 23, 2009

the connections from the → neurons on the left should be potentiated, and those from the right should be depressed, causing a leftward bias in the intracortical connections from these → neurons. For the same reason, a leftward moving object should cause a rightward bias in the connections from the → neurons. Thus, although in the natural environment right- and leftward motion stimuli are equally common, the above process operating during circuit development should lead to an asymmetry in the intracortical connections from direction-selective cells in the adult visual cortex. This asymmetry predicts two novel effects of motion stimuli on RF position. First, motion signals in stationary test stimuli should cause a displacement of the RF in the direction opposite to motion. Second, adaptation to motion stimuli should induce a shift of the RF in the direction of adaptation. Both predictions have been confirmed experimentally in adult cat visual cortex (41). Furthermore, comparison between these effects observed in cat V1 and human psychophysical measurements indicate that these RF properties can largely account for two previously known motion-related illusions.

The notion that during early development motion signals can shape the direction-selective response properties of visual neurons was tested by examining the effect of moving bar stimuli on the visual responses of developing tectal neurons. Engart et al. (30) showed that exposure of *Xenopus* tadpole to repetitive unidirectional moving bar stimuli resulted in a rapid modification of the tectal circuit so that the response of the tectal neuron to the moving bar was selectively enhanced at the conditioned direction. This modification was persistent after conditioning and involved NMDAR-dependent potentiation and depression of different subpopulations of retinotectal connections (122). By repetitively pairing light bar stimuli with the spiking of the tectal neuron, Mu et al. (76a) found that the magnitude and polarity of the changes in the light-evoked excitatory synaptic responses in tectal neurons exhibited a temporal specificity consistent with STDP shown in Figure 1. Furthermore, spike timing-dependent LTP and LTD of the retinotectal synapses have been shown to require BDNF and nitric oxide (NO) signaling, respectively, and the direction-selective enhancement of tectal responses induced by the moving bar stimuli was abolished by the inhibitor of trkB signaling and NO synthesis (76a). These findings all support the idea that STDP mediates experience-dependent circuit refinement in the developing neural circuits.

In addition to motion signals in the sensory stimuli, locomotion of the animals may also shape the nervous system through STDP. Repeated locomotion of the rat along a linear track induces an asymmetric expansion of the hippocampal place field (71), which can be accounted for by STDP of the hippocampal CA3-CA1 connections (1). This form of place field plasticity may underlie sequence learning during spatial navigation, which was predicted in a theoretical study (1). Because various types of motion are common in the natural environment, they are likely to constitute an important class of events that shape the neural circuits through STDP.

Finally, experience-dependent plasticity of the somatosensory cortical map may also involve STDP. Using a linear electrode array to record simultaneously from L4 and L2/3 neurons in the same cortical column, Celikel et al. (18) found that trimming the principle whisker corresponding to the recorded cortical column, such that it escapes the stimulation applied simultaneously to all whiskers, causes an immediate reversal of the order of spiking of L4 and L2/3 neurons, with only modest effects on the firing rate. According to the STDP window for these L4-L2/3 synapses measured in slices (31), the relative spike timing of L4 and L2/3 neurons during normal multiwhisker stimulation should cause little synaptic modification, while the whisker trimming-induced change in spike timing should cause significant LTD of the L4-L2/3 synapses. Indeed, depression of L4-L2/3 synaptic transmission has been demonstrated in barrel cortical slices from animals with whisker deprivation (5), and such trimming-induced depression occluded further LTD and enhanced LTP of these synapses in slices. These results suggest that STDP is also involved in the developmental cortical modifications induced by sensory deprivation.

D. Persistence of STDP In Vivo

Compared with in vitro preparations, neural circuits in vivo usually exhibit a higher level of spontaneous spiking, which may influence activity-dependent synaptic modifications. Recordings from freely moving adult rats showed that electrically induced LTP in the hippocampus is quickly reversed when the rat enters a novel environment within 1 h after induction (116). In developing retinotectal synapses, LTP and LTD induced by either electrical stimulation (paired spiking or theta bursts) or visual stimuli are rapidly reversed by spontaneous spiking of the tectal neuron or by random visual stimuli (122). This reversal of LTP can be prevented if postsynaptic spiking or NMDAR activation is blocked during the first 20 min after induction but not afterwards. The susceptibility of LTP and LTD to reversal by spontaneous activity in vivo may account for the short-term nature of STDP of cortical RFs and human visual perception (40, 117). In the retinotectal system, the reversal of LTP can be prevented by repeating the induction protocol a few times at an optimal spacing (122). Such a requirement for "spaced" induction protocol may prevent accidental pairing of pre/post spikes from causing lasting changes in the neural circuits.

Downloaded from physrev.physiology.org on October 23, 2009

The cellular mechanism underlying the reversal of LTP by spontaneous activity in vivo may be similar to that underlying the homosynaptic "de-potentiation" in hippocampal slices, in which low-frequency stimulation of the same pathway within tens of minutes following LTP induction protocol causes the reversal of LTP. In both cases the LTP reversal depends on the activity of protein phosphatases (see review in Ref. 121).

V. NONSYNAPTIC ASPECTS OF SPIKE TIMING-DEPENDENT PLASTICITY

A. STDP of Intrinsic Neuronal Excitability

Persistent changes in neuronal excitability following repetitive activity have long been noted in the studies of associative learning in various systems (see reviews in Refs. 21, 120). Changes in postsynaptic neuronal firing characteristics can result from modulation of either the excitatory/inhibitory synaptic drive or the intrinsic membrane conductances (intrinsic excitability). The latter has been reported in many systems. For example, activity associated with trace eye-blink conditioning in rabbits causes an enhanced intrinsic excitability of hippocampal CA1 and CA3 pyramidal cells for a few days, as shown by the reduction in both spike accommodation during prolonged depolarization and postburst afterhyperpolarization (76). In cerebellar deep nuclei, which are implicated in the trace eye-blink conditioning, brief high-frequency activation of mossy fiber inputs also causes a rapid and persistent increase in the intrinsic excitability (4). Similarly, theta burst stimulation (TBS) of mossy fibers potentiates the intrinsic excitability of granule cells in the cerebellum, an effect that can be dissociated from TBS-induced LTP of mossy fiber-granule cell synapses (9). Interestingly, although these changes in intrinsic excitability depend on NMDAR activation, they can occur in the absence of LTP, suggesting separate downstream mechanisms. The immediate changes of intrinsic excitability in all these studies point to the existence of rapid cytoplasmic signaling mechanisms that cause global modulation of ion channels in the neuronal membrane.

A form of spike timing-dependent modification of intrinsic neuronal excitability was also reported in hippocampal cultures, where the induction of LTP/LTD by correlated pre/post spiking is accompanied by an immediate and persistent enhancement/reduction of the intrinsic excitability of the presynaptic neuron, as reflected in a shift of the firing threshold and in the increased/decreased firing rate evoked by a constant depolarizing current injection at the soma (42, 58). This modification of excitability is temporally specific, with a requirement for pre/post spiking intervals identical to that for LTP/LTD (see Fig. 4). Similar to LTP/LTD induction, changes in the presynaptic excitability require activation of NMDARs and postsynaptic influx of Ca^{2+}, indicating that transsynaptic retrograde signaling is required. While presynaptic loading of PKC inhibitor (calphostin C or Ro-31–8220) through the whole cell pipette eliminated the spiking-induced increase in excitability, it only caused slight reduction of LTP. Thus, although the increased presynaptic excitability may contribute to LTP, it is not the primary cause (see below). Direct voltage-clamp measurements of Na^+ currents at the soma further showed that the increased excitability associated with LTP is related to changes in Na^+ channel activation and inactivation kinetics in favor of spike initiation. In the case of LTD, similar pharmacological and voltage-clamp studies showed that the reduction in excitability is related to an enhancement of the activation of slow-inactivating K^+ channels and that the activity of both protein kinase C (PKC) and protein kinase A (PKA) is required. Given the rapidity of these presynaptic changes (on the order of minutes) and the requirement for presynaptic PKA/PKC activities, post-translational modification of Na^+ and K^+ channels is the most likely underlying mechanism. Because the excitability measurements were carried out at the soma, it remains unclear to what extent the excitability is modified at the nerve terminal membrane, where changes in Na^+ and K^+ channel kinetics can affect the initiation and time course of the AP, which can in turn modulate evoked transmitter release and thus contribute to activity-induced changes in synaptic transmission.

In cell cultures, each presynaptic neuron makes many more synapses with the postsynaptic cell than that

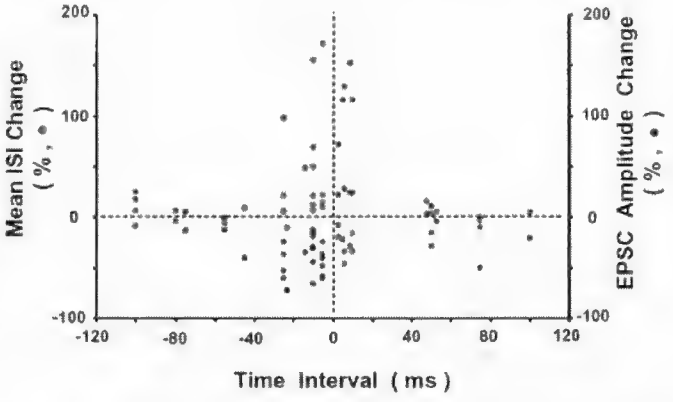

FIG. 4. Change in presynaptic excitability, measured by interspike interval (ISI) during depolarizing current injection, following the pre/post-spike pairing protocol in hippocampal culture. Each data point represents result from one experiment. The percentage changes for ISI (red) and excitatory postsynaptic current (EPSC) amplitude (blue) were calculated from the mean value at 10–20 min after correlated activation, compared with the mean value during the control period. The interval refers to the time between the onset of the excitatory postsynaptic potential and the peak of the postsynaptic spike.

Downloaded from physrev.physiology.org on October 23, 2009

in vivo, due to the artificial spatial proximity of the pre- and postsynaptic cells. Thus it is of interest to note that such global presynaptic modification was also observed in brain slices. Increased global excitability of the presynaptic cell was observed in L2/3 interneuron-pyramidal cell synapses in somatosensory cortical slices following the induction of spike timing-dependent LTD (58). Whole cell recording from cell pairs in which the presynaptic cell is a projection neuron with a long axon is much more difficult to achieve. It is possible that, in the latter situation, the retrograde influence due to LTP/LTD induction at the axonal terminal will be spatially more restricted, and global changes in excitability may not be detectable at the soma of the presynaptic cell.

B. STDP of Local Dendritic Excitability and Synaptic Integration

In addition to changes in the global presynaptic excitability, correlated activity also results in the modification of local postsynaptic excitability. In the original study of Bliss and Lomo (14), induction of LTP by tetanic stimulation was followed by an increase in the coupling between EPSPs and postsynaptic spiking ("E-S potentiation") that is distinct from the enhanced synaptic transmission (see Refs. 6, 20, 87). A reduction of the tonic inhibitory drive may contribute to E-S potentiation following LTP induction (3, 19, 61, 105), but local modification of dendritic conductances is also implicated (10, 45, 48). Recent dendritic recordings showed that LTP induction is indeed accompanied by a local modulation of transient A-type K^+ currents that can account for the enhanced excitability associated with E-S potentiation (36). In addition to LTP-associated changes, Daoudal et al. (22) found that LTD is also accompanied by an NMDAR-dependent, persistent E-S depression. This effect is input specific and is expressed when GABA receptors are blocked, suggesting a reduction in the intrinsic excitability of postsynaptic dendrites. In general, modulation of local active conductances depends on NMDAR activation and may represent a direct consequence of LTP/LTD induction. Local postsynaptic modulation of transmitter receptors, including their state of phosphorylation and the trafficking in and out of the subsynaptic membrane, have been implicated in LTP/LTD (11). Local modulation of voltage-dependent ion channels may be mediated by similar mechanisms.

Modulation of local active conductances in the dendrite can influence not only the initiation and propagation of dendritic spikes, but also the summation of synaptic potentials, a critical step in neuronal processing of information. In hippocampal CA1 pyramidal neurons, a brief period of correlated pre/post spiking that induces LTP and LTD also results in a persistent increase and decrease, respectively, in the linearity of spatial summation of EPSPs (111). These modifications are input specific, i.e., they occur only for the summation of the modified input with other inputs on the same postsynaptic dendrite. The increase in linearity was attributed primarily to a local modification of I_h channels, which are known to influence dendritic summation of EPSPs (17, 63, 64). These changes in linearity accompanying LTP/LTD help to boost the effects of synaptic modification by further enhancing/reducing the contribution of the modified input to the firing of the postsynaptic neuron.

VI. SPREAD OF SYNAPTIC SPIKE TIMING-DEPENDENT PLASTICITY IN NEURAL CIRCUITS

A. Heterosynaptic Effects of LTP/LTD Induction

Activity-induced synaptic modifications are often thought to be input specific, i.e., only the synapses experiencing repetitive synaptic activation are modified (7, 26, 62, 80). However, there is also substantial evidence for two forms of "breakdown" of input specificity. First, the induction of LTP or LTD at one input to the postsynaptic cell may cause opposite changes in the synaptic efficacy of other adjacent inputs to the same cell. A typical example is the heterosynaptic LTD found in the CA1 region of the hippocampus following the induction of LTP (2, 16, 27, 62, 90). A similar heterosynaptic effect also occurs in the amygdala (89) and the cortex (46). The second form of breakdown is related to the heterosynaptic effects of the same polarity following LTP/LTD induction, leading to a spread of potentiation/depression to adjacent synapses. In hippocampal CA1 pyramidal cells, LTP induced by pairing pyramidal cell depolarization with stimulation of the Schaffer collateral inputs was found to spread to synapses made by Schaffer collaterals on neighboring pyramidal cells (93). In the same cells, LTD induced by low-frequency Schaffer collateral stimulation (100) or by correlated post-pre spiking (81) also spreads to other inputs to the same cell or reverse LTP previously induced in a separate pathway (78). In the CA1 region of the hippocampus, induction of LTD of excitatory synapses onto GABAergic interneurons by tetanic stimulation of the input fibers leads to a spread of depression to the neighboring excitatory synapses onto the same interneuron (70).

Using an elegant local perfusion method that restricted synaptic activation to a small region (<20 μm) of the hippocampal slice culture, Engert and Bonhoeffer (29) have estimated the extent of spread of LTP in the CA1 pyramidal cell. They found that LTP induced at one

Downloaded from physrev.physiology.org on October 23, 2009

Schaffer collateral input to the pyramidal cell by high-frequency stimulation could spread to other synapses onto the same pyramidal cell when the unstimulated inputs were within ~70 μm from the site of LTP induction. The mechanism underlying this lateral spread of LTP remains unclear. It may be mediated by cytoplasmic signals in the postsynaptic cell or by membrane-permeant diffusible factors that spread through the extracellular space. Given the finding that LTP induced at one pyramidal cell can spread to synapses on adjacent pyramidal cells (93), it appears that local extracellular signals are involved in spreading the potentiation. The localized nature of these signals is also consistent with earlier findings of input specificity of LTP, where unstimulated control inputs were usually chosen to be relatively far from the stimulated input.

B. Spread of LTP/LTD in Cell Culture

Understanding the rules governing the nonlocal effects of LTP/LTD is essential for a full understanding of activity-induced modifications of neural circuits. Multiple whole cell recordings from small networks of cultured hippocampal neurons showed that LTP induction by correlated pre-post stimulation of a pair of connected glutamatergic neurons is accompanied by two forms of spread of potentiation (102): a "retrograde spread" to the inputs onto the dendrite of the presynaptic neuron and a "presynaptic lateral spread" of potentiation to synapses made by the presynaptic neuron onto other postsynaptic cells. The spread of potentiation in this culture system was specific in three aspects: *1)* there was no "forward (anterograde) spread" of potentiation to the output synapses of the postsynaptic neuron. *2*) There was no "postsynaptic lateral" spread of potentiation to converging inputs onto the postsynaptic neuron made by other presynaptic neurons, suggesting that input specificity of LTP is preserved. *3*) Presynaptic lateral spread was found only for glutamatergic terminals that innervate other glutamatergic neurons, but not GABAergic neurons.

How does one reconcile the finding of the restricted LTP spread in cultured hippocampal slices (29) with the absence of LTP spread in dissociated cell cultures described above? One possibility is that the average distance between synapses made by converging inputs in dissociated cell cultures exceed 70 μm, the extent of spread found in slice cultures. Alternatively, the disorganized distribution of synaptic sites in dissociated neurons may have reduced the magnitude and spatial distribution of the spread signal. The target cell-specific presynaptic lateral spread of LTP is reminiscent of the target cell-specific action of brain-derived neurotrophic factor (BDNF) found in these hippocampal cultures, where BDNF-induced potentiation of glutamatergic synaptic transmission was observed only for synapses made onto glutamatergic but not GABAergic neurons (91). It is possible that BDNF signaling associated with LTP induction results in a retrograde signaling, e.g., internalized endosomes containing active BDNF-trkB complexes, that spread to presynaptic nerve terminals onto other unstimulated postsynaptic cells. By having target cell specificity, correlated activity-induced LTP may still preserve some aspect of the input specificity. Given such presynaptic lateral spread of LTP, synaptic modification may be restricted to the set of presynaptic nerve terminals of the same axon that innervate the same type of postsynaptic target cells. Finally, it appears that the spread of potentiation is restricted entirely to the synapses received or made by the presynaptic neuron involved in LTP induction. This is reminiscent of the finding of global enhancement of excitability of the presynaptic neuron following LTP induction in these cultures (see above, Ref. 42), suggesting similar retrograde signaling may be involved in both phenomena.

For GABAergic synapses, lateral spread of synaptic modification was examined in cultured hippocampal neurons (115). No presynaptic spread of modification by correlated pre- and postsynaptic spiking was observed between divergent outputs of the same GABAergic neuron onto two different postsynaptic cells. However, for convergent GABAergic synapses onto the same postsynaptic neuron, substantial spread of modification to unstimulated inputs was observed in some cases. The variability in the extent of the spread may result from differences in the relative dendritic location of the convergent inputs.

Extensive retrograde and lateral spread of LTD has also been observed in cultures of hippocampal neurons, using the conventional protocol of pairing presynaptic stimulation with steady postsynaptic depolarization (32). Whether spike timing-dependent LTD exhibits similar spread remains to be determined. However, in acute hippocampal slices, Nishiyama et al. (81) showed that LTD induced by post-pre spiking at Schaffer collateral-CA1 pyramidal cell synapses spreads extensively to other converging inputs to the same postsynaptic pyramidal cell, although LTP induced at the same synapses by pre-post spiking is input specific. The induction of homo- and heterosynaptic LTD requires functional ryanodine receptors and inositol trisphosphate ($InsP_3$) receptors, respectively. Interestingly, pharmacological blockade or genetic deletion of type 1 $InsP_3$ receptors leads to a conversion of LTD to LTP and the elimination of heterosynaptic LTD, whereas blocking ryanodine receptors eliminates only homosynaptic LTD. These results suggest that postsynaptic Ca^{2+} elevation derived from both Ca^{2+} influx and differential release of Ca^{2+} from internal stores through ryanodine and $InsP_3$ receptors can regulate both the po-

Downloaded from physrev.physiology.org on October 23, 2009

larity and the spread of spike timing-dependent synaptic modification. An $InsP_3$ receptor-dependent Ca^{2+} wave in the postsynaptic neuron may underlie the spread of Ca^{2+} elevation throughout the postsynaptic dendrite, leading to the spread of LTD.

C. Spread of LTP/LTD In Vivo

The findings on the spread of LTP/LTD in cultured cells or slice preparations raise the question of whether similar spreads occur in vivo. Recent studies using intact retinotectal system of developing *Xenopus* tadpoles have addressed this issue. In this system, LTP induced by correlated pre-post spiking of the retinal ganglion cell (RGC) and tectal neuron was found to be input specific throughout early development; there was no spread of potentiation laterally to other converging RGC inputs on the postsynaptic tectal cell or to synapses made by the same RGC axons to other adjacent tectal cells (103). However, LTP of the same retinotectal synapses induced by conventional TBS of RGCs was found to spread extensively to other inputs during the early stages of development, and input specificity of this LTP emerged as the tectal cell develops more elaborate dendritic arbors (103). This difference in the input specificity of LTP induced by correlated spiking and TBS during early development points to an important distinction of STDP. Compared with the conventional forms of LTP, STDP may code for a more specific form of synaptic modification that requires a more precise pattern of spiking activity. Fluorescence Ca^{2+} imaging showed that a global Ca^{2+} elevation in these young tectal neurons was triggered by TBS but not by repetitive correlated spiking (103), suggesting that the input specificity may be attributed to a more localized postsynaptic Ca^{2+} elevation. The emergence of input specificity appears to correlate with the development of spatially restricted domains of Ca^{2+} elevation in the tectal cell dendrite triggered by each retinal axon input, consistent with the idea that Ca^{2+} signaling may participate in the spread of LTP.

The most intriguing observation on the spread of LTP/LTD in hippocampal cultures is the rapid retrograde (or "back-propagation") of potentiation/depression to the synaptic inputs onto the dendrite of the presynaptic neuron. Recent study in the retinotectal system suggests that extensive back-propagation of synaptic modification may also exist in vivo. Persistent potentiation of retinotectal synapses by application of BDNF at the tectum is accompanied by a potentiation of synaptic inputs to the dendrite of the RGCs, leading to increased light-evoked RGC response (25). The BDNF effect at the dendrite depends on TrkB expression in the RGCs and is absent when the optic nerve is severed. Further analysis showed that a rapid increase of AMPA receptors at the RGC dendrites was induced by the exposure of retinotectal synapses to BDNF. This retrograde synaptic modification is similar to the back-propagation of LTP found in cultured hippocampal neurons, suggesting retrograde transport of signals via uptake of factors at the nerve terminal may provide a mechanism for the spread of LTP/LTD in the presynaptic neuron. The same signal may also account for the bidirectional modification of global intrinsic excitability of the presynaptic neuron (see above).

Recent studies have further shown that LTP/LTD of *Xenopus* retinotectal synapses induced by TBS/low-frequency stimulation of the optic nerve is accompanied by a long-range retrograde spread of potentiation/depression from the RGC axon terminals to the synaptic input onto the RGC dendrites. Such retrograde potentiation and depression could be abolished by severing the optic nerve or by blocking the retrograde signaling mediated by BDNF and NO, respectively. Whether spike timing-dependent LTP/LTD also undergo similar retrograde spread remains to be determined. The existence of such long-range retrograde spread of LTP/LTD allows adjustment of the strength of input synapses at the dendrite of a neuron in accordance with that of the output synapses at its axon terminals. Such a mechanism may be useful for a coordinated circuit refinement, in which activity-dependent potentiation/depression of the outputs of a neuron will be used by the circuit to carry out corresponding changes in its input signals, resulting in the consolidation of correct pathways.

VII. CONCLUDING REMARKS

While much progress has been made over the past decade in understanding the role of spike timing in long-term synaptic modification, research in this area is still at its infancy. Most of our current knowledge of STDP is based on studies of excitatory synapses on pyramidal cells. Given the diversity of cell types and their distinct biophysical properties, there is likely to be a variety of cell type specific STDP windows, which remain to be characterized. How synaptic and nonsynaptic forms of STDP operate in concert to shape the functions of neural circuits is only beginning to be explored. Furthermore, sensory processing and many other behavioral tasks involve sequences of events separated by seconds to minutes. How STDP at the time scale of tens of milliseconds relates to these integrative brain functions is largely unknown. Finally, while this review focuses mainly on experimental studies, computational analysis of the consequences of STDP, especially for circuits involving large populations of neurons, will be critical for our understanding of the function of STDP from synapse to perception.

Downloaded from physrev.physiology.org on October 23, 2009

ACKNOWLEDGMENTS

Address for reprint requests and other correspondence: Y. Dan (e-mail: ydan@berkeley.edu) or M. Poo (e-mail: mpoo@berkeley.edu), Div. of Neurobiology, Dept. of Molecular and Cell Biology, and Helen Wills Neuroscience Institute, Univ. of California, Berkeley, CA 94720.

REFERENCES

1. **Abbott LF and Blum KI.** Functional significance of long-term potentiation for sequence learning and prediction. *Cereb Cortex* 6: 406–416, 1996.
2. **Abraham WC and Goddard GV.** Asymmetric relationships between homosynaptic long-term potentiation and heterosynaptic long-term depression. *Nature* 305: 717–719, 1983.
3. **Abraham WC, Gustafsson B, and Wigstrom H.** Long-term potentiation involves enhanced synaptic excitation relative to synaptic inhibition in guinea-pig hippocampus. *J Physiol* 394: 367–380, 1987.
4. **Aizenman CD and Linden DJ.** Rapid, synaptically driven increases in the intrinsic excitability of cerebellar deep nuclear neurons. *Nat Neurosci* 3: 109–111, 2000.
5. **Allen CB, Celikel T, and Feldman DE.** Long-term depression induced by sensory deprivation during cortical map plasticity in vivo. *Nat Neurosci* 6: 291–299, 2003.
6. **Andersen P, Sundberg SH, Sveen O, Swann JW, and Wigstrom H.** Possible mechanisms for long-lasting potentiation of synaptic transmission in hippocampal slices from guinea-pigs. *J Physiol* 302: 463–482, 1980.
7. **Andersen P, Sundberg SH, Sveen O, and Wigstrom H.** Specific long-lasting potentiation of synaptic transmission in hippocampal slices. *Nature* 266: 736–737, 1977.
8. **Archie KA and Mel BW.** A model for intradendritic computation of binocular disparity. *Nat Neurosci* 3: 54–63, 2000.
9. **Armano S, Rossi P, Taglietti V, and D'Angelo E.** Long-term potentiation of intrinsic excitability at the mossy fiber-granule cell synapse of rat cerebellum. *J Neurosci* 20: 5208–5216, 2000.
10. **Asztely F and Gustafsson B.** Dissociation between long-term potentiation and associated changes in field EPSP waveform in the hippocampal CA1 region: an in vitro study in guinea pig brain slices. *Hippocampus* 4: 148–156, 1994.
11. **Bear MF and Malenka RC.** Synaptic plasticity: LTP and LTD. *Curr Opin Neurobiol* 4: 389–399, 1994.
12. **Bell CC, Han VZ, Sugawara Y, and Grant K.** Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387: 278–281, 1997.
13. **Bi GQ and Poo MM.** Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* 18: 10464–10472, 1998.
14. **Bliss TV and Lomo T.** Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J Physiol* 232: 331–356, 1973.
15. **Boettiger CA and Doupe AJ.** Developmentally restricted synaptic plasticity in a songbird nucleus required for song learning. *Neuron* 31: 809–818, 2001.
16. **Bradler JE and Barrioneuvo G.** Long-term potentiation in hippocampal CA3 neurons: tetanized input regulates heterosynaptic efficacy. *Synapse* 4: 132–142, 1989.
17. **Cash S and Yuste R.** Linear summation of excitatory inputs by CA1 pyramidal neurons. *Neuron* 22: 383–394, 1999.
18. **Celikel T, Szostak VA, and Feldman DE.** Modulation of spike timing by sensory deprivation during induction of cortical map plasticity. *Nat Neurosci* 7: 534–541, 2004.
19. **Chavez-Noriega LE, Bliss TV, and Halliwell JV.** The EPSP-spike (E-S) component of long-term potentiation in the rat hippocampal slice is modulated by GABAergic but not cholinergic mechanisms. *Neurosci Lett* 104: 58–64, 1989.
20. **Chavez-Noriega LE, Halliwell JV, and Bliss TV.** A decrease in firing threshold observed after induction of the EPSP-spike (E-S) component of long-term potentiation in rat hippocampal slices. *Exp Brain Res* 79: 633–641, 1990.
21. **Daoudal G and Debanne D.** Long-term plasticity of intrinsic excitability: learning rules and mechanisms. *Learn Mem* 10: 456–465, 2003.
22. **Daoudal G, Hanada Y, and Debanne D.** Bidirectional plasticity of excitatory postsynaptic potential (EPSP)-spike coupling in CA1 hippocampal pyramidal neurons. *Proc Natl Acad Sci USA* 99: 14512–14517, 2002.
23. **Debanne D, Gahwiler BH, and Thompson SM.** Cooperative interactions in the induction of long-term potentiation and depression of synaptic excitation between hippocampal CA3-CA1 cell pairs in vitro. *Proc Natl Acad Sci USA* 93: 11225–11230, 1996.
24. **Debanne D, Gahwiler BH, and Thompson SM.** Long-term synaptic plasticity between pairs of individual CA3 pyramidal cells in rat hippocampal slice cultures. *J Physiol* 507: 237–247, 1998.
25. **Du JL and Poo MM.** Rapid BDNF-induced retrograde synaptic modification in a developing retinotectal system. *Nature* 429: 878–883, 2004.
26. **Dudek SM and Bear MF.** Homosynaptic long-term depression in area CA1 of hippocampus and effects of *N*-methyl-D-aspartate receptor blockade. *Proc Natl Acad Sci USA* 89: 4363–4367, 1992.
27. **Dunwiddie T and Lynch G.** Long-term potentiation and depression of synaptic responses in the rat hippocampus: localization and frequency dependency. *J Physiol* 276: 353–367, 1978.
28. **Egger V, Feldmeyer D, and Sakmann B.** Coincidence detection and changes of synaptic efficacy in spiny stellate neurons in rat barrel cortex. *Nat Neurosci* 2: 1098–1105, 1999.
29. **Engert F and Bonhoeffer T.** Synapse specificity of long-term potentiation breaks down at short distances. *Nature* 388: 279–284, 1997.
30. **Engert F, Tao HW, Zhang LI, and Poo MM.** Moving visual stimuli rapidly induce direction sensitivity of developing tectal neurons. *Nature* 419: 470–475, 2002.
31. **Feldman DE.** Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27: 45–56, 2000.
32. **Fitzsimonds RM, Song HJ, and Poo MM.** Propagation of activity-dependent synaptic depression in simple neural networks. *Nature* 388: 439–448, 1997.
33. **Fregnac Y, Shulz D, Thorpe S, and Bienenstock E.** A cellular analogue of visual cortical plasticity. *Nature* 333: 367–370, 1988.
34. **Fregnac Y and Shulz DE.** Activity-dependent regulation of receptive field properties of cat area 17 by supervised Hebbian learning. *J Neurobiol* 41: 69–82, 1999.
35. **Freund TF and Buzsaki G.** Interneurons of the hippocampus. *Hippocampus* 6: 347–470, 1996.
36. **Frick A, Magee J, and Johnston D.** LTP is accompanied by an enhanced local excitability of pyramidal neuron dendrites. *Nat Neurosci* 7: 126–135, 2004.
37. **Froemke RC and Dan Y.** Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* 416: 433–438, 2002.
38. **Froemke RC, Poo MM, and Dan Y.** Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434: 221–225, 2005.
39. **Froemke RC, Tsay IA, Raad M, Long JD, and Dan Y.** Contribution of individual spikes in burst-induced long-term synaptic modification. *J Neurophysiol.* In press.
40. **Fu YX, Djupsund K, Gao H, Hayden B, Shen K, and Dan Y.** Temporal specificity in the cortical plasticity of visual space representation. *Science* 296: 1999–2003, 2002.
41. **Fu YX, Shen Y, Gao H, and Dan Y.** Asymmetry in visual cortical circuits underlying motion-induced perceptual mislocalization. *J Neurosci* 24: 2165–2171, 2004.
42. **Ganguly K, Kiss L, and Poo M.** Enhancement of presynaptic neuronal excitability by correlated presynaptic and postsynaptic spiking. *Nat Neurosci* 3: 1018–1026, 2000.
43. **Golding NL, Staff NP, and Spruston N.** Dendritic spikes as a mechanism for cooperative long-term potentiation. *Nature* 418: 326–331, 2002.
44. **Hausser M and Mel B.** Dendrites: bug or feature? *Curr Opin Neurobiol* 13: 372–383, 2003.

Downloaded from physrev.physiology.org on October 23, 2009

45. **Hess G and Gustafsson B.** Changes in field excitatory postsynaptic potential shape induced by tetanization in the CA1 region of the guinea-pig hippocampal slice. *Neuroscience* 37: 61–69, 1990.
46. **Hirsch JC, Barrionuevo G, and Crepel F.** Homo- and heterosynaptic changes in efficacy are expressed in prefrontal neurons: an in vitro study in the rat. *Synapse* 12: 82–85, 1992.
47. **Huber KM, Kayser MS, and Bear MF.** Role for rapid dendritic protein synthesis in hippocampal mGluR-dependent long-term depression. *Science* 288: 1254–1257, 2000.
48. **Jester JM, Campbell LW, and Sejnowski TJ.** Associative EPSP-spike potentiation induced by pairing orthodromic and antidromic stimulation in rat hippocampal slices. *J Physiol* 484: 689–705, 1995.
49. **Johnston D, Christie BR, Frick A, Gray R, Hoffman DA, Schexnayder LK, Watanabe S, and Yuan LL.** Active dendrites, potassium channels and synaptic plasticity. *Philos Trans R Soc Lond B Biol Sci* 358: 667–674, 2003.
50. **Kampa BM, Clements J, Jonas P, and Stuart GJ.** Kinetics of Mg^{2+} unblock of NMDA receptors: implications for spike-timing dependent synaptic plasticity. *J Physiol* 556: 337–345, 2004.
51. **Karmarkar UR and Buonomano DV.** A model of spike-timing dependent plasticity: one or two coincidence detectors? *J Neurophysiol* 88: 507–513, 2002.
52. **Kepecs A, van Rossum MC, Song S, and Tegner J.** Spike-timing-dependent plasticity: common themes and divergent vistas. *Biol Cybern* 87: 446–458, 2002.
53. **Kirkwood A and Bear MF.** Homosynaptic long-term depression in the visual cortex. *J Neurosci* 14: 3404–3412, 1994.
54. **Klausberger T, Magill PJ, Marton LF, Roberts JD, Cobden PM, Buzsaki G, and Somogyi P.** Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* 421: 844–848, 2003.
55. **Kobayashi K and Poo MM.** Spike train timing-dependent associative modification of hippocampal CA3 recurrent synapses by mossy fibers. *Neuron* 41: 445–454, 2004.
56. **Koester HJ and Sakmann B.** Calcium dynamics in single spines during coincident pre- and postsynaptic activity depend on relative timing of back-propagating action potentials and subthreshold excitatory postsynaptic potentials. *Proc Natl Acad Sci USA* 95: 9596–9601, 1998.
57. **Levy WB and Steward O.** Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8: 791–797, 1983.
58. **Li C, Lu J, Wu C, Duan S, and Poo M.** Bidirectional modification of presynaptic neuronal excitability accompanying spike timing-dependent synaptic plasticity. *Neuron* 41: 257–268, 2004.
59. **Lin YW, Min MY, Chiu TH, and Yang HW.** Enhancement of associative long-term potentiation by activation of beta-adrenergic receptors at CA1 synapses in rat hippocampal slices. *J Neurosci* 23: 4173–4181, 2003.
60. **Liu QS, Pu L, and Poo M.** Repeated cocaine exposure in vivo facilitates LTP induction in midbrain dopamine neurons. *Nature*. In press.
61. **Lu YM, Mansuy IM, Kandel ER, and Roder J.** Calcineurin-mediated LTD of GABAergic inhibition underlies the increased excitability of CA1 neurons associated with LTP. *Neuron* 26: 197–205, 2000.
62. **Lynch GS, Dunwiddie T, and Gribkoff V.** Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature* 266: 737–739, 1977.
63. **Magee JC.** Dendritic hyperpolarization-activated currents modify the integrative properties of hippocampal CA1 pyramidal neurons. *J Neurosci* 18: 7613–7624, 1998.
64. **Magee JC.** Dendritic Ih normalizes temporal summation in hippocampal CA1 neurons. *Nat Neurosci* 2: 848, 1999.
65. **Magee JC and Johnston D.** A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275: 209–213, 1997.
66. **Malenka RC and Siegelbaum SA.** Synaptic plasicity. In: *Synapses*, edited by Cowan WM, Sudhof TC, and Stevens CF. Baltimore, MD: Johns Hopkins Univ. Press, 2001.
67. **Markram H, Lubke J, Frotscher M, and Sakmann B.** Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275: 213–215, 1997.
68. **Martin SJ, Grimwood PD, and Morris RG.** Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu Rev Neurosci* 23: 649–711, 2000.
69. **Mayer ML, Westbrook GL, and Guthrie PB.** Voltage-dependent block by Mg^{2+} of NMDA responses in spinal cord neurones. *Nature* 309: 261–263, 1984.
70. **McMahon LL and Kauer JA.** Hippocampal interneurons express a novel form of synaptic plasticity. *Neuron* 18: 295–305, 1997.
71. **Mehta MR, Quirk MC, and Wilson MA.** Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25: 707–715, 2000.
72. **Mel BW, Ruderman DL, and Archie KA.** Translation-invariant orientation tuning in visual "complex" cells could derive from intradendritic computations. *J Neurosci* 18: 4325–4334, 1998.
73. **Meliza CD and Dan Y.** Receptive-field modification in rat visual cortex induced by paired visual stimulation and single cell spiking. *Neuron*. In press.
74. **Meredith RM, Floyer-Lea AM, and Paulsen O.** Maturation of long-term potentiation induction rules in rodent hippocampus: role of GABAergic inhibition. *J Neurosci* 23: 11142–11146, 2003.
75. **Montgomery JM and Madison DV.** State-dependent heterogeneity in synaptic depression between pyramidal cell pairs. *Neuron* 33: 765–777, 2002.
76. **Moyer JR Jr, Thompson LT, and Disterhoft JF.** Trace eyeblink conditioning increases CA1 excitability in a transient and learning-specific manner. *J Neurosci* 16: 5536–5546, 1996.
76a. **Mu Y and Poo MM.** Spike-timing-dependent LTP/LTD mediates visual experience-dependent plasticity in a developing retinotectal system. *Neuron* 50: 115–125, 2006.
77. **Mulkey RM and Malenka RC.** Mechanisms underlying induction of homosynaptic long-term depression in area CA1 of the hippocampus. *Neuron* 9: 967–975, 1992.
78. **Muller D, Hefft S, and Figurov A.** Heterosynaptic interactions between LTP and LTD in CA1 hippocampal slices. *Neuron* 14: 599–605, 1995.
79. **Nevian T and Sakmann B.** Single spine Ca^{2+} signals evoked by coincident EPSPs and backpropagating action potentials in spiny stellate cells of layer 4 in the juvenile rat somatosensory barrel cortex. *J Neurosci* 24: 1689–1699, 2004.
80. **Nicoll RA and Malenka RC.** Neurobiology. Long-distance long-term depression. *Nature* 388: 427–428, 1997.
81. **Nishiyama M, Hong K, Mikoshiba K, Poo MM, and Kato K.** Calcium stores regulate the polarity and input specificity of synaptic modification. *Nature* 408: 584–588, 2000.
82. **Normann C, Peckys D, Schulze CH, Walden J, Jonas P, and Bischofberger J.** Associative long-term depression in the hippocampus is dependent on postsynaptic N-type Ca^{2+} channels. *J Neurosci* 20: 8290–8297, 2000.
83. **Nowak L, Bregestovski P, Ascher P, Herbet A, and Prochiantz A.** Magnesium gates glutamate-activated channels in mouse central neurones. *Nature* 307: 462–465, 1984.
84. **Oliet SH, Malenka RC, and Nicoll RA.** Two distinct forms of long-term depression coexist in CA1 hippocampal pyramidal cells. *Neuron* 18: 969–982, 1997.
85. **Polsky A, Mel BW, and Schiller J.** Computational subunits in thin dendrites of pyramidal cells. *Nat Neurosci* 7: 621–627, 2004.
86. **Pouille F and Scanziani M.** Routing of spike series by dynamic circuits in the hippocampus. *Nature* 429: 717–723, 2004.
87. **Pugliese AM, Ballerini L, Passani MB, and Corradetti R.** EPSP-spike potentiation during primed burst-induced long-term potentiation in the CA1 region of rat hippocampal slices. *Neuroscience* 62: 1021–1032, 1994.
88. **Rosenmund C, Feltz A, and Westbrook GL.** Calcium-dependent inactivation of synaptic NMDA receptors in hippocampal neurons. *J Neurophysiol* 73: 427–430, 1995.
89. **Royer S and Pare D.** Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature* 422: 518–522, 2003.

Downloaded from physrev.physiology.org on October 23, 2009

90. **Scanziani M, Malenka RC, and Nicoll RA.** Role of intercellular interactions in heterosynaptic long-term depression. *Nature* 380: 446–450, 1996.
91. **Schinder AF, Berninger B, and Poo M.** Postsynaptic target specificity of neurotrophin-induced presynaptic potentiation. *Neuron* 25: 151–163, 2000.
92. **Schuett S, Bonhoeffer T, and Hubener M.** Pairing-induced changes of orientation maps in cat visual cortex. *Neuron* 32: 325–337, 2001.
93. **Schuman EM and Madison DV.** Locally distributed synaptic potentiation in the hippocampus. *Science* 263: 532–536, 1994.
94. **Segev I and London M.** Untangling dendrites with quantitative models. *Science* 290: 744–750, 2000.
95. **Shouval HZ, Bear MF, and Cooper LN.** A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc Natl Acad Sci USA* 99: 10831–10836, 2002.
96. **Sjostrom PJ, Turrigiano GG, and Nelson SB.** Neocortical LTD via coincident activation of presynaptic NMDA and cannabinoid receptors. *Neuron* 39: 641–654, 2003.
97. **Sjostrom PJ, Turrigiano GG, and Nelson SB.** Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* 32: 1149–1164, 2001.
98. **Softky WR and Koch C.** The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13: 334–350, 1993.
99. **Song S, Miller KD, and Abbott LF.** Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3: 919–926, 2000.
100. **Staubli UV and Ji ZX.** The induction of homo- vs. heterosynaptic LTD in area CA1 of hippocampal slices from adult rats. *Brain Res* 714: 169–176, 1996.
101. **Stefan K, Kunesch E, Benecke R, Cohen LG, and Classen J.** Mechanisms of enhancement of human motor cortex excitability induced by interventional paired associative stimulation. *J Physiol* 543: 699–708, 2002.
102. **Tao H, Zhang LI, Bi G, and Poo M.** Selective presynaptic propagation of long-term potentiation in defined neural networks. *J Neurosci* 20: 3233–3243, 2000.
103. **Tao HW, Zhang LI, Engert F, and Poo M.** Emergence of input specificity of ltp during development of retinotectal connections in vivo. *Neuron* 31: 569–580, 2001.
104. **Togashi K, Kitajima T, Aihara T, Hong K, Poo M, and Nishiyama M.** Gating of activity-dependent long-term depression by GABAergic activity in the hippocampus. *Soc Neurosci Abstr* 123.4, 2003.
105. **Tomasulo RA and Ramirez JJ.** Activity-mediated changes in feed-forward inhibition in the dentate commissural pathway: relationship to EPSP/spike dissociation in the converging perforant path. *J Neurophysiol* 69: 165–173, 1993.
106. **Tong G, Shepherd D, and Jahr CE.** Synaptic desensitization of NMDA receptors by calcineurin. *Science* 267: 1510–1512, 1995.
107. **Tzounopoulos T, Kim Y, Oertel D, and Trussell LO.** Cell-specific, spike timing-dependent plasticities in the dorsal cochlear nucleus. *Nat Neurosci* 7: 719–725, 2004.
108. **Umemiya M, Chen N, Raymond LA, and Murphy TH.** A calcium-dependent feedback mechanism participates in shaping single NMDA miniature EPSCs. *J Neurosci* 21: 1–9, 2001.
109. **Vicini S, Wang JF, Li JH, Zhu WJ, Wang YH, Luo JH, Wolfe BB, and Grayson DR.** Functional and pharmacological differences between recombinant *N*-methyl-D-aspartate receptors. *J Neurophysiol* 79: 555–566, 1998.
110. **Wang HX, Gerkin RC, Nauen DW, and Bi GQ.** Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat Neurosci* 8: 187–193, 2005.
111. **Wang Z, Xu NL, Wu CP, Duan S, and Poo MM.** Bidirectional changes in spatial dendritic integration accompanying long-term synaptic modifications. *Neuron* 37: 463–472, 2003.
112. **Watanabe S, Hoffman DA, Migliore M, and Johnston D.** Dendritic K^+ channels contribute to spike-timing dependent long-term potentiation in hippocampal pyramidal neurons. *Proc Natl Acad Sci USA* 99: 8366–8371, 2002.
113. **Wigstrom H and Gustafsson B.** Postsynaptic control of hippocampal long-term potentiation. *J Physiol* 81: 228–236, 1986.
114. **Wolters A, Sandbrink F, Schlottmann A, Kunesch E, Stefan K, Cohen LG, Benecke R, and Classen J.** A temporally asymmetric Hebbian rule governing plasticity in the human motor cortex. *J Neurophysiol* 89: 2339–2345, 2003.
115. **Woodin MA, Ganguly K, and Poo MM.** Coincident pre- and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in Cl^- transporter activity. *Neuron* 39: 807–820, 2003.
116. **Xu L, Anwyl R, and Rowan MJ.** Spatial exploration induces a persistent reversal of long-term potentiation in rat hippocampus. *Nature* 394: 891–894, 1998.
117. **Yao H and Dan Y.** Stimulus timing-dependent plasticity in cortical processing of orientation. *Neuron* 32: 315–323, 2001.
118. **Yao H, Shen Y, and Dan Y.** Intracortical mechanism of stimulus-timing-dependent plasticity in visual cortical orientation tuning. *Proc Natl Acad Sci USA* 101: 5081–5086, 2004.
119. **Zhang LI, Tao HW, Holt CE, Harris WA, and Poo M.** A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395: 37–44, 1998.
120. **Zhang W and Linden DJ.** The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat Rev Neurosci* 4: 885–900, 2003.
121. **Zhou Q and Poo MM.** Reversal and consolidation of activity-induced synaptic modifications. *Trends Neurosci* 27: 378–383, 2004.
122. **Zhou Q, Tao HW, and Poo MM.** Reversal and stabilization of synaptic modifications in a developing visual system. *Science* 300: 1953–1957, 2003.

Downloaded from physrev.physiology.org on October 23, 2009

nature neuroscience

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

The tempotron: a neuron that learns spike timing–based decisions

Robert Gütig[1–4] & Haim Sompolinsky[1,2,5]

The timing of action potentials in sensory neurons contains substantial information about the eliciting stimuli. Although the computational advantages of spike timing–based neuronal codes have long been recognized, it is unclear whether, and if so how, neurons can learn to read out such representations. We propose a new, biologically plausible supervised synaptic learning rule that enables neurons to efficiently learn a broad range of decision rules, even when information is embedded in the spatiotemporal structure of spike patterns rather than in mean firing rates. The number of categorizations of random spatiotemporal patterns that a neuron can implement is several times larger than the number of its synapses. The underlying nonlinear temporal computation allows neurons to access information beyond single-neuron statistics and to discriminate between inputs on the basis of multineuronal spike statistics. Our work demonstrates the high capacity of neural systems to learn to decode information embedded in distributed patterns of spike synchrony.

Uncovering the coding principles by which neurons in the CNS represent and process information has challenged brain researchers for over half a century. It is commonly assumed that neurons represent information through their mean rates of action potential firing. However, in visual, auditory, olfactory and somatosensory pathways[1–5], spikes elicited in response to a stimulus can be precisely timed relative to a stimulus event or to other action potentials of the same or other neurons. These data suggest that the temporal structure of spike trains serves as an important component of the neuronal representation of stimuli. For instance, in the human tactile system[5], the latencies of action potentials in individual tactile afferents carry information about the direction of the external force and the shape of the contacting surface. Spike latency codes have also been suggested for the olfactory system[6] and retinal ganglion cells[7]. The synchrony between spike trains of neuronal pairs in visual cortex has been implicated in the representation of stimulus 'gestalt' features[8]. Multineuronal synchronous events have been detected in the retina and have receptive fields different from those predicted by the firing rates of the individual neurons[1] (but see ref. 9). Indeed, some researchers have argued that, under certain circumstances, temporal neural codes offer significant computational advantages over rate codes[6,7,10–12]. Others have argued, however, that temporal codes are inferior because of the complexity of their decoding[9,13].

These experimental and computational issues highlight two important and long-standing questions. First, through what mechanisms of synaptic plasticity can neurons learn to discriminate between different spatiotemporal sequences of incoming spike patterns? That is, how do neurons learn to read out information carried by a temporal code? Second, what spatiotemporal features embedded within incoming spike patterns do neurons use to compute their response? Most existing computational models of supervised learning, such as the powerful perceptron[14,15] (**Supplementary Note** online), are formulated in a rate-based framework and do not apply to the learning of spike timing–based decision rules. This shortcoming has obstructed theoretical and experimental advances in understanding the role of spike-timing in neural information processing and learning.

To close this gap, we have devised a new, biologically plausible model of supervised learning, the tempotron, for decoding information embedded in spatiotemporal spike patterns. Using tempotron learning, we show that an integrate-and-fire neuron can learn to categorize a broad range of input classes. These include cases in which category information is not encoded in spike counts but, rather, in the latencies of single spikes or in pairwise and higher-order patterns of synchrony.

RESULTS

The tempotron learning rule

Our neuron model consists of a leaky integrate-and-fire neuron driven by exponentially decaying synaptic currents generated by its N synaptic afferents. The subthreshold membrane voltage is a weighted sum of postsynaptic potentials (PSPs) from all incoming spikes:

$$V(t) = \sum_i \omega_i \sum_{t_i} K(t - t_i) + V_{\text{rest}} \qquad (1)$$

Here t_i denotes the spike times of the ith afferent and $K(t - t_i)$ is the normalized PSP contributed by each incoming spike: that is, $K(t - t_i) = V_0 (\exp[-(t - t_i)/\tau] - \exp[-(t - t_i)/\tau_s])$. The parameters τ and τ_s denote decay time constants of membrane integration and synaptic currents, respectively (**Fig. 1**). The factor V_0 normalizes PSP kernels to 1 so that

[1]Racah Institute of Physics and [2]Interdisciplinary Center for Neural Computation, Hebrew University, 91904 Jerusalem, Israel. [3]Institute for Theoretical Biology, Humboldt University, 10115 Berlin, Germany. [4]Neuroscience Research Center, Charitè Medical Faculty of Berlin, 10117 Berlin, Germany. [5]Center for Brain Science, Harvard University, Cambridge, Massachusetts 02138, USA. Correspondence should be addressed to H.S. (haim@fiz.huji.ac.il) or R.G. (guetig@cc.huji.ac.il).

Received 1 December 2005; accepted 17 January 2006; published online 12 February 2006; corrected 14 February 2006 (details online); doi:10.1038/nn1643

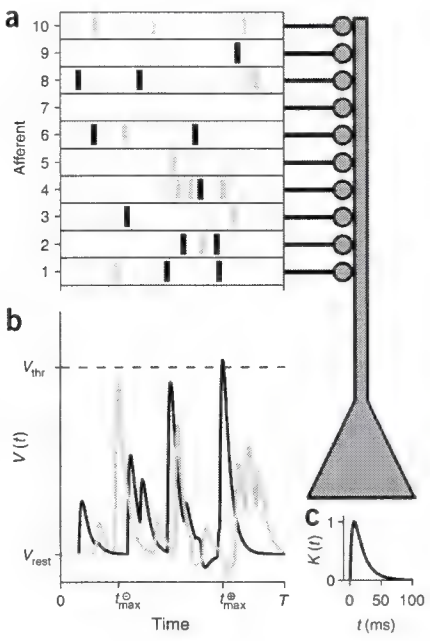

Figure 1 Tempotron classification. (**a**) Spike times (thick vertical lines) of ten afferents for two schematic input patterns, ⊕ (black) and ⊖ (gray). (**b**) Resulting postsynaptic voltage traces $V(t)$. Maximal voltages were reached at $t_{max}^{\oplus}$ and $t_{max}^{\ominus}$, respectively. Because $V(t_{max}^{\oplus}) > V_{thr} > V(t_{max}^{\ominus})$ with V_{thr} (dashed horizontal line) denoting the spike threshold, both patterns are classified correctly. Inputs arriving after a threshold crossing (black trace) are shunted. (**c**) Postsynaptic integration kernel $K(t)$ with membrane time constant $\tau = 15$ ms and synaptic time $\tau_s = \tau/4$. Here and in all other figures, the input arrival times are between 0 and $T = 500$ ms.

unitary PSP amplitudes are given by the synaptic efficacies ω_i. $K(t - t_i)$ is a causal filter, vanishing for $t_i > t$. When $V(t)$ crossed the firing threshold, the neuron is said to elicit a spike (referred to as an output spike), and the voltage undergoes a smooth reset to V_{rest} by shunting all incoming spikes that arrived after the output spike (Methods).

In the classification tasks considered here, each input pattern belongs to one of two classes (which are labeled ⊕ and ⊖). A pattern consists of N trains of input spikes, one train for each afferent, arriving between times 0 and T. The tempotron's task is to respond to a ⊕ pattern by firing at least one action potential and to remain quiescent when driven by a ⊖ pattern (**Fig. 1**). The tempotron learns this task by a rule that modifies the synaptic efficacies ω_i whenever an error occurs. In an iterative learning scheme, synapses should be depressed by an amount that reflects their contribution to an erroneous output spike on a ⊖ pattern or potentiated according to their responsibility for failure to spike on a ⊕ pattern. To resolve the temporal credit assignment problem of gauging the relative contribution of afferents to an erroneous decision, the tempotron uses the following rule: If no output spike was elicited in response to a ⊕ pattern, each synaptic efficacy ω_i is increased by the following amount:

$$\Delta\omega_i = \lambda \sum_{t_i < t_{max}} K(t_{max} - t_i) \qquad (2)$$

Here t_{max} denotes the time at which the postsynaptic potential $V(t)$ reaches its maximal value. The constant $\lambda > 0$ specifies the maximum size of the synaptic update per input spike. Conversely, if an output spike appears in response to a ⊖ pattern the synaptic efficacies are decreased by $\Delta\omega_i$ (**Fig. 2**).

The tempotron learning rule implements a 'gradient-descent' dynamics, which minimizes a cost function that measures the amount by which the maximum voltage generated by erroneous patterns deviates from the firing threshold. For each ⊕ pattern that did not elicit a spike, the cost function is $V_{thr} - V(t_{max})$. For each ⊖ pattern that induced an output spike, the cost function is $V(t_{max}) - V_{thr}$. In gradient-based learning, the change in the synaptic efficacy is proportional to the negative of the derivative of the cost function with respect to ω_i. Differentiating the value of the voltage at the peak with respect to the efficacies (Methods) gives the tempotron learning rule, shown in equation (2). Using this rule, we studied the ability of a leaky integrate-and-fire neuron to learn the classification of spatiotemporal spike patterns that mimic scenarios of spike latency (**Fig. 3a**) and spike synchrony codes (**Fig. 3b,c**).

Learning to classify latency patterns

In our first test of the tempotron's ability to recognize spatiotemporal spike features, we assessed its performance in classifying populations of single spike latencies. The task consisted of p spike patterns, each of which was randomly assigned to one of the two classes ⊕ or ⊖. In each pattern, each afferent fired once at a fixed time chosen randomly from a uniform distribution between 0 and $T = 500$ ms (**Fig. 3a**). The ability of the tempotron to classify such random patterns correctly depends crucially on the richness of the filtering of its spatio-temporal synaptic integration, which in turn is a function of the number of synapses. Therefore, a useful measure of the load on the system is the number of patterns p relative to the number of synapses, denoted by $\alpha = p/N$. An important characteristic of the neuron's capacity is the maximal load that it can learn. The tempotron can learn to perform the task as long as α is less than a critical value, approximately 3 (**Fig. 4a**). The similarity of the dependence of the learning time on α for different values of N (**Fig. 4a**) indicates that the maximum number of patterns that the tempotron can learn grows linearly with the number of afferents. Notably, this capacity exceeds the known capacity of $\alpha = 2$ of a single-layer perceptron[14–16] (**Supplementary Note** and **Supplementary Fig. 1** online), reflecting the higher complexity of the computation performed by the integrate-and-fire neuron across time. In the classification task studied above, the spike count was the same for all inputs in all patterns; hence, the pattern's class membership was specified entirely by the spike latencies. By contrast, an example of rate coding could be implemented through spike patterns in each of which a randomly chosen half of the afferents do not spike while the other half fire synchronously[17] (Methods). In this case, the tempotron learning rule reduces to the perceptron learning rule (see also **Fig. 4a**, left curve). Thus, notably, the tempotron can read both spike timing–based and rate-based codes.

The computation performed by the neuron depends on the PSP integration window relative to the stimulus duration. For concreteness, we varied τ and τ_s while holding the ratio between them fixed ($\tau = 4\tau_s$). For τ much shorter than the average time between input spikes T/N, no temporal summation of separated inputs takes place. In contrast, slow kinetics renders all spikes effectively simultaneous. Hence, in both limits, the extraction of temporal features is impaired and differentiation between ⊕ and ⊖ patterns on basis of spike timing becomes impossible. Correspondingly, the tempotron solves the random latency classification task within an intermediate range of PSP integration times (**Fig. 4b**). This range and the optimal τ, for which learning was fastest, depend on the number of synapses N. Increasing N substantially reduces the lower limit of τ (data not shown). Additionally, whereas the optimal τ is ~10 ms for $N = 500$, it reduces to ~4 ms

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

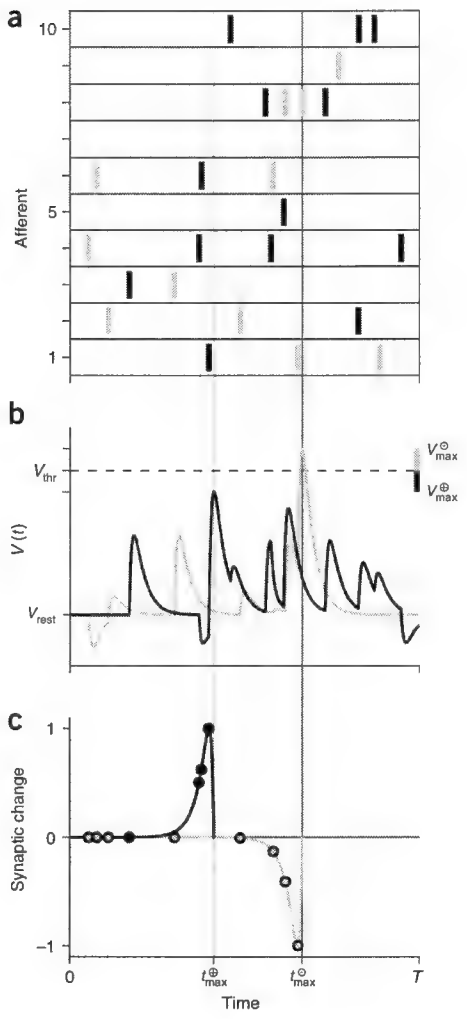

Figure 2 Tempotron learning rule. (**a**) Spike times (thick vertical lines) of ten afferents in two schematic input patterns, ⊕ (black) and ⊖ (gray). (**b**) Resulting postsynaptic voltage traces $V(t)$. Maximal voltages are reached at $t_{max}^{\oplus}$ and $t_{max}^{\ominus}$ (thin solid vertical lines), respectively. Because $V(t_{max}^{\oplus}) < V_{thr} < V(t_{max}^{\ominus})$ with V_{thr} (dashed horizontal line) denoting the spike threshold, both patterns generate an error. Inputs arriving after a threshold crossing (gray trace) are shunted. Black and gray thick vertical lines indicate the cost terms $V_{max} - V_{thr}$ and $V_{thr} - V_{max}$ associated with the ⊕ and ⊖ patterns, respectively. (**c**) Resulting synaptic changes (in units of maximal change, λ) depend on presynaptic spike times (circles) relative to the corresponding voltage maximum. Thin dashed vertical lines in (**a–c**) mark presynaptic spike times.

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

for $N = 1{,}000$ (see also **Fig. 4a**). The range of permissible values of τ shrinks with increasing load α (**Fig. 4b**); however, for intermediate load, the range of integration time constants for which the tempotron solves the task is broad and covers sparse and dense modes of decision making. These modes are distinguished by the size of the input populations that contributed to a given decision to fire, N_{dec} (**Fig. 4c**, Methods). The mean of N_{dec} was 4.7 for $\tau = 2$ ms and 97 for $\tau = 75$ ms, confirming that the tempotron can operate in both sparse and dense regimes. The two cases differ also in the distributions of learned synaptic efficacies. Decisions based on integration of a small number of inputs require a substantial fraction of strong synaptic efficacies (**Fig. 4d**).

A characteristic feature of the tempotron is that the time of the output spike is not constrained by the classification task, allowing the neuron to distribute its decision times to optimize its performance. In the present task, the output spike times are broadly distributed: early output spikes appear within τ of the first input spikes and late output spikes appear at the end of the input patterns (**Fig. 5a**). This temporal homogeneity is also reflected in the mean subthreshold voltage traces, averaged separately over all ⊕ and ⊖ patterns after learning converged to zero error. After an onset ramp of duration τ, the mean voltage traces maintain a steady plateau value across the entire input time interval (**Fig. 5b**).

To assess the robustness of tempotron learning in the presence of temporal noise, we trained the neuron with jittered versions of p templates of latency patterns. In each presentation, a template was jittered by adding a Gaussian noise to all spike times. The tempotron's generalization error—that is, the probability of misclassifying a jittered spike pattern after the learning converged—grows with increasing jitter and higher learning load α (**Fig. 6a**). For a broad range of postsynaptic integration times τ, from roughly 2 ms to 75 ms, the tempotron can tolerate a jitter of roughly 1% of τ for $\alpha = 2$ and 10% for $\alpha = 1$ (data shown for $\tau = 15$ ms). Hence, with PSP kinetics on the order of tens of milliseconds and spike precisions in the millisecond range[5,11], the categorization of 1–2 templates per synapse is feasible. Notably, when discriminating between only two templates, the tempotron can overcome jitter of roughly 10τ. This demonstrates that for small load, the tempotron is able to learn temporal features that extend well beyond the timescale of its intrinsic dynamics. Another feature of the tempotron's robustness is the relative stability of the timing of output spikes to the temporal jitter of a ⊕ template. A small jitter in input spikes generates a distribution of output spikes for each template that has a prominent narrow peak around the unperturbed spike time. Input jitter causes a large change in the output spikes in only a few percent of the trials (**Fig. 6b,c**). The tempotron is also robust to the deletion or insertion of a small number of spikes. For $\alpha = 1$, its generalization error roughly equals the percentage of the perturbed spikes (data not shown).

Learning to classify patterns of synchrony

What information about the timing of spikes in ⊕ and ⊖ patterns does the tempotron imprint on its learned synaptic efficacies? Differences between the voltage traces averaged separately over the ⊕ and ⊖ random latency patterns are small relative to the fluctuations in individual traces (**Fig. 5b**), implying that the tempotron also relies on higher-order spike time statistics to separate the two input classes. To directly test the sensitivity of the tempotron to multineural firing statistics, we trained it to discriminate between classes of input patterns that cannot be distinguished on the basis of single neuron spike statistics. Specifically, we examined the simple scenario in which each input pattern consists of pairwise-synchronous events. All N afferents are grouped into $N/2$ pairs. Afferents in a given pair fire synchronously, each a single spike. Notably, this grouping is fixed for each of the two categories (⊕ and ⊖) but differs between them (**Figs. 3b** and **7a**; Methods). Synchronous events occur at random, uniformly distributed times in both pattern categories, so that class information is embedded solely in the patterns of synchrony; neither spike counts nor spike timing of individual neurons carry any information relevant for the classification task. This task mimicked spike synchrony–based sensory processing[4,8,18–20]. When trained on these patterns with a jitter of up to 2–3 ms, the tempotron quickly achieved practically zero generalization error. Learning splits the synapses into two groups, roughly equal in size, of excitatory ($\omega > 0$) and inhibitory ($\omega < 0$) efficacies (**Fig. 7b**).

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

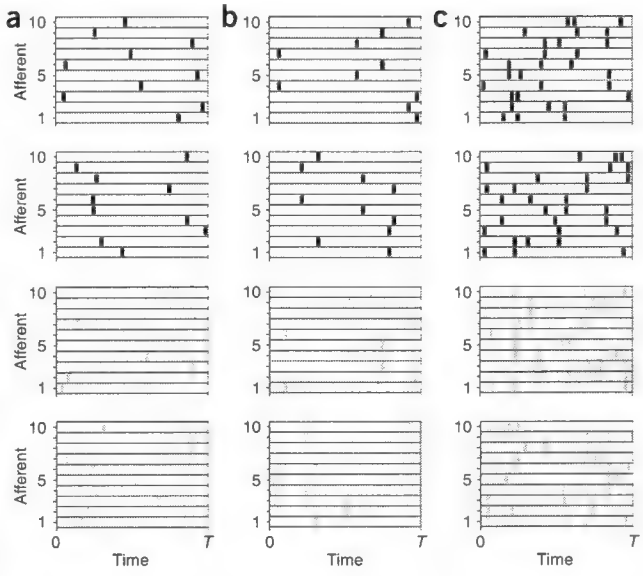

Figure 3 Input patterns (see Methods for details). Spike times (thick vertical lines) of ten afferents in input patterns representative of the three classification tasks considered in this work. The two upper rows depict two examples of ⊕ patterns (black rasters) and the two lower rows show two ⊖ patterns (gray rasters). (**a**) Random latencies. Each pattern consists of one spike per afferent, the time of which is drawn from a uniform distribution. In this classification task, class membership is assigned at random to the latency patterns. (**b**) Pairwise synchrony. Afferents in each input class, ⊕ and ⊖, are divided into $N/2$ synchronous pairs: for example, (1, 3) and (2, 10) in ⊕ patterns and (1, 6) and (2, 7) in ⊖ patterns. The identity of the synchronous pairs is conserved across all ⊕ and across all ⊖ patterns. In each realization of inputs, the arrival times of synchronous spikes are drawn from a uniform distribution. (**c**) Third-order interaction. In both input classes afferents are divided into the same groups: for example, (1, 2, 3), (4, 5, 6), (7, 8, 9), (alternating background shading). Each afferent fires three times during the input interval. In ⊕ patterns, each afferent participates in two synchronous spike pairs with its interacting neighbors and contributes one additional asynchronous spike. In contrast, in ⊖ patterns, each afferent participates in one synchronous spike triplet and contributes two additional single asynchronous spikes. In each realization, the times of all single, pair and triple events are random.

Pairs of synapses that are synchronized in ⊕ patterns have the same sign whereas pairs coincident in the ⊖ patterns have opposite signs. As a result, coincident inputs cancel on the ⊖ patterns but add in the ⊕ patterns. The nature of this solution reveals that the tempotron does not rely on a single synchronous pair but on a large number of pairs, using this redundancy to overcome ambiguities generated by the input jitter.

The ability of the tempotron to extract spike correlations is not limited to pairwise statistics. Motivated by indications of the potential importance of higher-order multineural spike statistics[21,22], we used synchronous spike pairs versus spike triplets to generate input classes that have the same first- and second-order statistics but differ in their third-order interactions. We did this by dividing all N afferents into $N/3$ groups of three. This grouping is conserved across all patterns and both classes. Depending on a pattern's class membership, the spike times within each such group of three afferents contain either three pairwise coincident events or a single coincident spike triplet (**Fig. 3c**; Methods). The tempotron can learn to discriminate these classes either by selectively firing in response to patterns with spike triplets (data not shown) or, even when required to detect the 'missing spikes', by firing exclusively in response to patterns with only pairwise coincidences (**Fig. 7c**). In contrast to the previous task (**Fig. 7b**), here the tempotron weakens almost all synapses and, by spontaneous symmetry breaking, selects a single group of three interacting afferents with which to compute its decision (**Fig. 7d**). Two of these weights are positive and one is negative. The magnitudes of the positive weights suffice to elicit an output spike when activated in synchrony in a ⊕ pattern but cannot overcome the negative contribution of the third coincident spike in the triplet in the ⊖ patterns. The strategy of computing the decision using a single group of interacting afferents minimizes the probability of errors due to random spike coincidences from other triplets. The ability to separate patterns according to higher-order spike time statistics reflects the underlying nonlinear computation across time performed by the tempotron, which goes beyond extensions of perceptron learning to the time domain[23].

Implementation by voltage convolution

One challenge for the biological implementation of our model is that the tempotron's solution to the temporal credit assignment problem requires computation of the time of the maximum total PSP. Notably,

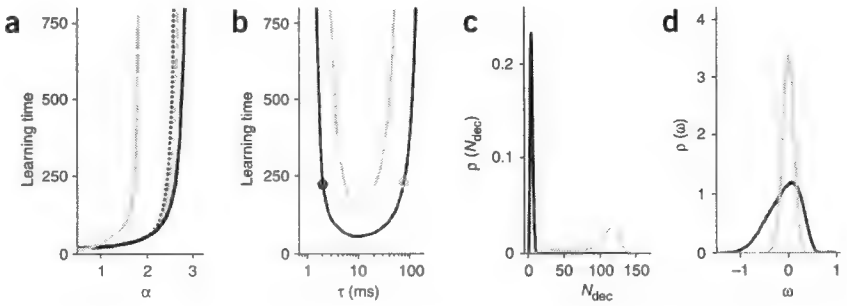

Figure 4 Tempotron performance. (**a**) Learning time—that is, mean number of cycles of pattern presentations required for error-free classification—versus load α. The load is defined as the ratio of the number of patterns over the number of afferents. Results for random latency patterns with $N = 500$ afferents (black solid line) were obtained with a synaptic integration time of $\tau = 10$ ms. For $N = 1{,}000$ afferents (gray dashed line) and $N = 1{,}500$ afferents (black dotted line) τ was set to 4 ms and 3 ms, respectively. These values of τ are close to the respective optimal τ. For comparison, the learning times for $N = 1{,}000$ and $N = 1{,}500$ were scaled to match the learning time for $N = 500$ at $\alpha = 2$. The gray solid line depicts learning times for perceptron-like input patterns with $N = 500$ afferents (Methods). (**b–d**) In **b**, learning time for $N = 500$ afferents versus PSP time constant τ for $\alpha = 2$ (black line) and $\alpha = 2.5$ (gray line). Open circles mark $\tau = 2$ ms (black) and $\tau = 75$ ms (gray) used in **c–d**. In **c**, distributions of effective number of contributing synapses N_{dec} (Methods) after learning ($\alpha = 2$) for $\tau = 2$ ms (black) and $\tau = 75$ ms (gray). In **d**, distributions of learned synaptic efficacies ω (in units of the spike threshold V_{thr}) for parameters as in **c** (same colors). Data are averaged over 1,000 realizations (Methods).

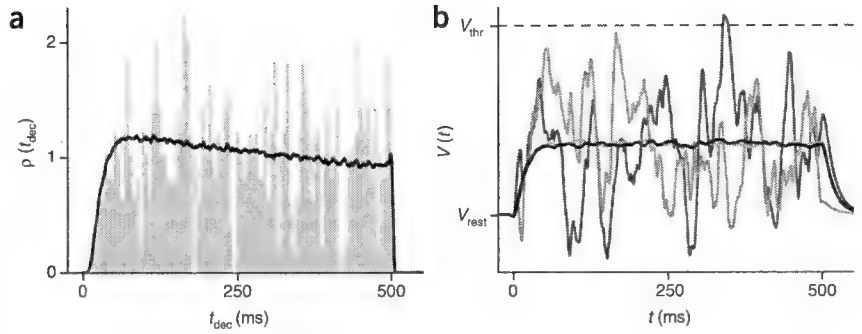

Figure 5 Tempotron decision times. (**a**) Distribution of output spike times on ⊕ patterns after learning ($\alpha = 2$, $N = 500$, $\tau = 15$ ms) for a single batch of $p = 1{,}000$ patterns (gray histogram) and average over 1,000 realizations (black line). (**b**) Typical voltage traces for individual ⊕ (blue) and ⊖ (red) patterns after learning and average voltages over all ⊕ (green) and ⊖ patterns (black) of one batch. Parameters as in **a**. The dashed horizontal line denotes the spike threshold V_{thr}. For the purpose of this comparison, we show the total PSP traces for the ⊕ patterns, ignoring the shunting effect of the output spike.

this temporally global operation can be approximated by temporally local computations that are based on the postsynaptic voltage traces following input spikes. Several functional forms that capture the temporal correlation between presynaptic spiking and postsynaptic voltage trace are possible. In one form of such a rule, the changes in synaptic efficacies following error trials are based on the convolution of the PSP kernel with the postsynaptic potential:

$$\upsilon_i = \int_{t_i}^{\infty} \mathrm{d}t\, V(t) K(t - t_i) \qquad (3)$$

The voltage maximum operation is approximated by thresholding υ_i: plasticity is induced at synapses for which υ_i is greater than a plasticity induction threshold κ. Efficacies are incremented by an amount λ for ⊕ patterns and by $-\lambda$ for ⊖ patterns (Methods). This rule ensures that synapses are changed only if the voltage within a postsynaptic integration time after their activation is sufficiently high. We found that, on all tasks considered here, this rule performs reasonably well compared to the tempotron rule. Although maximum capacity is somewhat reduced, by less than a factor of 2 (data not shown), for intermediate loads the learning time of the voltage convolution–based tempotron is only moderately increased (**Fig. 8**, inset). To investigate the relationship between this rule and the original tempotron rule, we computed the changes of all synapses that were activated at a given temporal offset from the time of maximal V, averaged separately for errors following ⊕ and ⊖ patterns. The resultant mean plasticity temporal profile (**Fig. 8**, black solid curve) is similar to that specified by the tempotron plasticity rule (**Fig. 8**, black dotted curve). This suggests that a temporally local mechanism that extracts the contiguity between the elevation of postsynaptic membrane potential and the timing of incoming spikes suffices to roughly mimic the tempotron's global synaptic credit assignment. This was also supported by our finding that voltage convolution learning converges to the same type of solution as that of tempotron learning in tasks involving multineuronal spike statistics (data not shown).

Comparison with reinforcement learning

Another approach to the credit assignment problem is to correlate the synaptic changes with neuronal noise, a method used by reinforcement learning. We implemented a synaptic reinforcement learning scheme[24] in which synaptic changes are governed by correlations between random synaptic noise, which was added to the synaptic efficacies, and a global error signal (Methods). Due to its inherent stochasticity, learning in this reinforcement scheme is dramatically slower than that in the tempotron (roughly four orders of magnitude in the example of **Fig. 8**, inset) and does not converge to zero error. We computed the mean synaptic changes induced by this learning rule as a function of the time of synaptic activation relative to the maximum $V(t)$. During the early phase of learning, the resultant temporal profile is substantially different from that of the tempotron (data not shown); in contrast, in the late phase both profiles are remarkably similar (**Fig. 8**, gray solid line) even though plasticity in the reinforcement learning scheme does not contain explicit reference to the postsynaptic voltage. This suggests that the computational principle underlying tempotron learning has a fundamental role in spike-based processing.

DISCUSSION

The tempotron model posits that synaptic changes depend on the value of the postsynaptic voltage. A crucial role of the postsynaptic voltage in the induction of long-term plasticity has been demonstrated in several

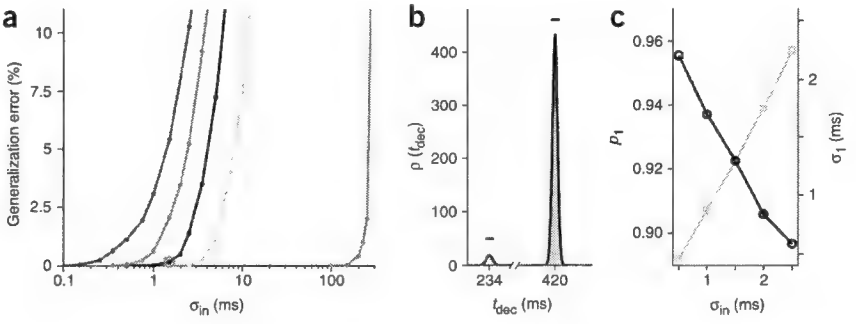

Figure 6 Tempotron robustness. (**a**,**b**) In **a**, generalization error—that is, the probability of misclassifying a jittered spike pattern after the learning converged—versus input jitter standard deviation σ_{in} for $\alpha = 2$ (blue), 1.5 (red), 1 (black), 0.5 (gray) and 1/250 (green). Results were obtained for $N = 500$ afferents and a synaptic time constant of $\tau = 15$ ms. Red open circle on $\alpha = 1$ curve (black) marks the σ_{in} used in **b**. In **b**, distribution of output times (gray histogram) originating from a single ⊕ pattern template with jitter $\sigma_{in} = 1.5$ ms after learning ($\alpha = 1$). The black line depicts the estimated density of a two-component Gaussian-mixture model (**Supplementary Methods**). Short horizontal scale bars centered above the mean of each component extend over $\pm\sigma_{in}$. (**c**) Mean weight p_1 (black circles, left *y*-axis) and mean standard deviation σ_1 (gray squares, right *y*-axis) of the dominant component of the distribution of output times versus σ_{in}. The quantities were computed by averaging the parameters of Gaussian-mixture fits across all ⊕ patterns. Data in **a** and **c** were averaged over five realizations.

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

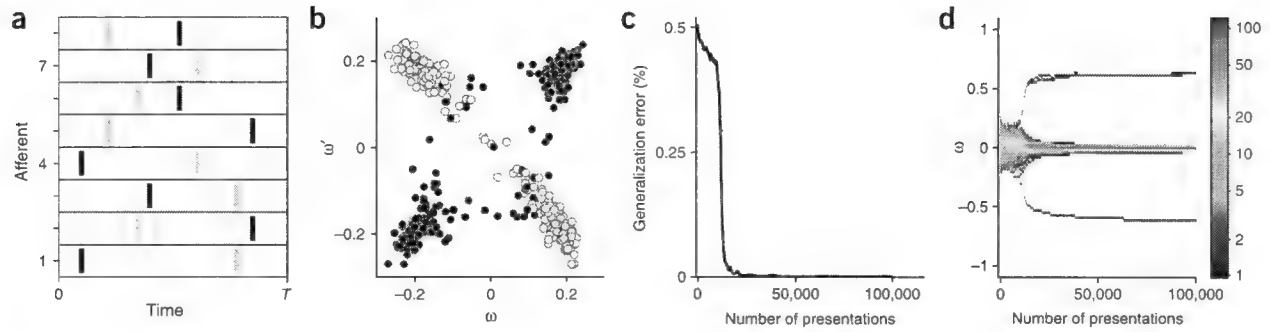

Figure 7 Learning multineural spike statistics. (**a**,**b**) Pairwise synchrony (see **Fig. 3b**). In **a**, spike times of eight afferents in a ⊕ (black) and a ⊖ (gray) pattern, realizing distinct pairings of synchronous input spikes. In **b**, scatter plot of pairs (ω, ω') of learned efficacies. Black circles show the efficacies grouped into the pairs that are active synchronously in ⊕ patterns: for example, $\omega = \omega_1$ and $\omega' = \omega_4$ (see **a**). Gray circles show the same efficacies now grouped into the pairs that are active synchronously in ⊖ patterns: for example, $\omega = \omega_7$ and $\omega' = \omega_4$ (see **a**). ($N = 500$, $\tau = 15$ ms). (**c**,**d**) Third-order interactions (**Fig. 3c**). In **c**, generalization error during learning of discrimination between patterns with different third-order interactions as a function of number of pattern presentations ($N = 168$, $\tau = 15$ ms). In **d**, time evolution of the histogram of synaptic efficacies, starting from a Gaussian distribution. Color-coded histograms of synaptic efficacies (logarithmic color scale, right bar) versus number of pattern presentations. Learning suppressed all but the three weights of a single group of interacting afferents that developed large magnitudes. Two of these were positive efficacies (upper branch) and the other was negative (lower branch). Sharp drop in generalization error in **c** accompanies the emergence of this solution.

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

preparations[25–28]. In addition, local voltage computations similar to those discussed above can be implemented by the thresholding properties of voltage-gated calcium channels and NMDA receptors, whose involvement in the induction of long-term plasticity is well established[25,29,30]. However, the voltage in the tempotron rule is the postsynaptic potential, which in the case of an erroneous output spike (on a ⊖ pattern) does not incorporate the voltage produced by the active conductances. The tempotron incorporates the effect of opening the active conductances by erasing all the synaptic inputs that arrive after an output spike.

Most experimentally derived plasticity rules, such as spike timing–dependent plasticity (STDP), characterize synaptic changes solely in terms of the temporal contiguity of the presynaptic spike and the postsynaptic potential or spiking. Indeed, several computational studies have shown that such 'unsupervised' rules can extract salient statistical features in the input spike trains[31–36]. In contrast, the tempotron solves the more challenging problem of learning to respond to presynaptic spike sequences by making the appropriate decision. Such learning necessarily requires supervisory signals, which instruct the neuron about the correct decision. An important question concerning the biological feasibility of tempotron learning is through which pathways supervisory signals arrive at the site of plasticity and how they are translated into the appropriate synaptic changes.

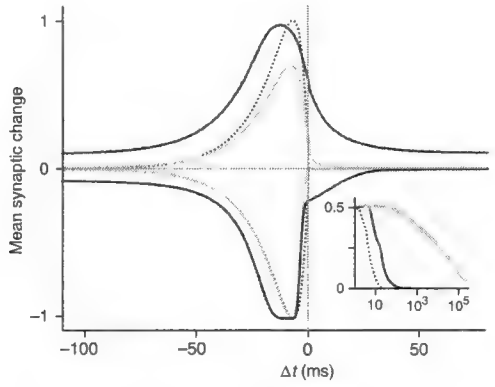

Figure 8 Spike timing dependence of different learning rules. Mean synaptic changes versus the time difference Δt between presynaptic spikes and postsynaptic voltage maximum, for voltage convolution–based (solid black), reinforcement (solid gray) and tempotron learning (dotted black, equation (2)). Minima are scaled to −1. Positive (negative) curves correspond to mean synaptic changes following errors on ⊕ (⊖) patterns. Data were collected during the late phase of learning, after the fraction of misclassified patterns dropped below 0.25. The learning task consisted of random latency patterns (**Fig. 3a**) with $N = 500$ afferents and a load of $\alpha = 1$ ($\tau = 15$ ms). The inset (same color code and parameter values) shows the fraction of misclassified patterns versus the number of cycles of pattern presentations during learning. All data were averaged over 1,000 realizations (solid and dotted black) and 10 realizations (gray).

In the tempotron, the magnitudes of synaptic changes are fully determined by the time course of the postsynaptic voltage. The polarity of synaptic changes is the same for all synapses and only depends on the desired response. Hence, the role of the supervisory signal is to determine this global polarity. One way to implement the control of polarity is through classic error feedback. In such a scenario, an error trial is followed either by a signal that results in long-term potentiation (LTP) if the neuron failed to spike on a ⊕ pattern or by a signal resulting in long-term depression (LTD) if the neuron erroneously fired on a ⊖ pattern. The minimal requirements for such a scheme are twofold. First, synapses must be able to maintain voltage-dependent 'eligibility' traces over substantial periods of time, such that the subsequent delivery of the error feedback can be translated into the correct amount of change. Second, there must be neuromodulatory signaling pathways that can be recruited to activate either LTP or LTD after an error trial, depending on the type of error. One problematic aspect of this feedback scenario is that it relies on the calculation and transmission of the error by the supervisory system, which may considerably delay the execution of the appropriate synaptic modification. This sequence of calculations can be simplified by realizing that the tempotron requires only the availability of a signal that instructs the desired response and not an explicit computation of error. We can build on this feature to construct a scheme in which a synaptic depression process is induced independently of the desired response whenever the tempotron fires an output spike. In contrast, potentiation does not depend on postsynaptic spiking but is induced only in the presence of a neuromodulator released by the supervisory system. This activation of the LTP process is the point of interaction through which

the tempotron is instructed that the desired response is to spike. Neuromodulatory control of LTP occurs for several neuromodulatory pathways, such as dopamine and acetylcholine[37–39]. Finally, it is necessary that coactivation of both processes results in cancellation of potentiation and depression, such that no changes in synaptic efficacies are expressed if the neuron spikes correctly in response to a ⊕ pattern. On a ⊖ pattern, no plasticity is induced when the neuron remains quiescent, as neither LTP nor LTD is elicited. However, the failure to spike on ⊕ patterns results in potentiation whereas spiking on ⊖ patterns results in depression, producing the correct polarity of synaptic changes on erroneous ⊕ and ⊖ trials.

Recent reports provide support for the long-standing hypothesis[40,41] that the polarity of expressed synaptic change results from the balance between separately activated intracellular induction pathways for LTP and LTD (refs. 42,43), both of which converge onto the trafficking of postsynaptic AMPA receptors[44–46]. Due to its intricate dependence on a large number of intracellular state and signaling variables, this balance between a depression process modulated by calcineurin[41] and a potentiation process mediated by Ca^{2+}/calmodulin-dependent protein kinase II (CaMKII; ref. 47) offers a large number of possible interaction sites for both neuromodulatory and state-dependent regulations[37,41,48]. At present, these intracellular signaling cascades are not yet fully exposed, precluding the determination of the feasibility of any detailed scenario of balancing between LTP and LTD processes. Independent of the details of how tempotron learning is implemented, our work predicts that, unlike simple STDP rules, some synaptic plasticity in behaving animals is governed by two factors: the temporal contiguity of a presynaptic spike and a postsynaptic depolarization, and a global neuromodulatory signal that determines the polarity of the expressed long-term synaptic changes.

Our results highlight the ability of simple neuronal circuits to learn and implement spike timing–based computation. This supports the robust decoding of a broad range of spatiotemporal neural codes, in contrast to earlier assertions[9,13]. A natural question is what are the computational limitations of the tempotron. In the tempotron, output spikes are elicited in response to spatiotemporal input features whose duration is on the order of the postsynaptic integration window (τ in our notation). As shown here, the tempotron can learn to detect the presence of such temporally restricted features anywhere in a stimulus time interval (T in our notation), even when this interval is significantly larger than τ. However, the tempotron is expected to perform poorly on classification tasks that require it to discriminate between spike patterns on the basis of temporal features that extend beyond a single integration time[49]. For instance, this limitation restricts the tempotron's ability to learn the conjunction of temporally local features separated by long times. Solving such demanding tasks requires the addition of working memory mechanisms or slow synaptic dynamics[50] and possibly also multilayer architectures.

METHODS

Neuron model. Numerical simulations of the tempotron were based on exact integration of the voltage dynamics of a leaky integrate-and-fire neuron driven by exponentially decaying synaptic currents (see equation (1)). For concreteness, we held the ratio between the membrane and the synaptic time constants (τ and τ_s, respectively) fixed at $\tau/\tau_s = 4$. With this ratio, the normalization coefficient V_0 (see equation (1)) equaled 2.12. The important temporal parameter is the ratio τ/T. We chose $T = 500$ ms and $\tau = 15$ ms (except in **Fig. 4**). PSP amplitudes of individual afferents were given by the synaptic efficacies ω_i. In this model, the sign of the synaptic efficacy could change by learning. This can be realized by altering the balance between excitatory and inhibitory pathways. We checked that the tempotron learning rule also worked (although with a reduced capacity) when the signs of the synaptic efficacies were not allowed to change. The neuron was simulated for up to 10,000 plastic synapses; results shown are for $N = 500$ synapses, except for capacity measurements which are also shown for $N = 1{,}000$ and $N = 1{,}500$ synapses (**Fig. 4a**) and for the learning of higher-order spike statistics where three spikes arrived per afferent and $N = 168$ (**Fig. 7c,d**).

An output spike was elicited when $V(t)$ crossed the firing threshold V_{thr}. After a spike at t_{spike}, the voltage was reset smoothly to the resting value by shunting all synaptic inputs that arrived after the occurrence of the spike—namely, by restricting the sum in equation (1) to $t_i < t_{spike}$ (see Tempotron Learning, below). We used $V_{thr} = 1$ and $V_{rest} = 0$.

Input patterns. As stated above, for all patterns, the duration of the input time interval [0, T] was set to $T = 500$ ms. In random patterns (**Figs. 3a, 4–6** and **8**), each afferent spiked once at a time drawn independently from a uniform distribution. Perceptron-like patterns (**Fig. 4a**, gray curve) consisted of a random half of afferents spiking in synchrony while the other half remained silent. Patterns were assigned to input classes ⊕ or ⊖ with a probability of one-half. Jittered realizations of patterns (**Fig. 6**) were generated from the original patterns by adding independent Gaussian noise with mean = 0 and standard deviation σ_{in} to all spike arrival times while keeping the class identity of the original pattern. In learning patterns of pairwise synchrony (**Fig. 7a,b**), afferents were randomly grouped into $N/2$ pairs, the identities of which were different for each of the two classes ⊕ and ⊖ (**Fig. 3b**). In each pattern, each pair of afferents contributed one synchronous spike pair at a time drawn randomly from a uniform distribution. In learning patterns with higher-order spike correlations (**Fig. 7c,d**), afferents were assigned into groups of three, with identical and fixed identities for both pattern classes (**Fig. 3c**). Each afferent spiked three times during the input interval. In ⊕ patterns, two of the spikes of each group member formed synchronous spike pairs with its two neighbors. In ⊖ patterns, one spike of each group member engaged in a triplet of synchronized spikes. In both classes, the remaining spikes did not interact. To prevent overlap between events within each interacting group of three afferents, spike times were drawn from a uniform distribution but separated by a distance of $\tau + \tau_s$.

Tempotron learning. Synapses were modified after each error trial. We defined the single pattern cost functions $E_\pm = \pm (V_{thr} - V(t_{max})) \Theta(\pm (V_{thr} - V(t_{max}))$ for ⊕ and ⊖ patterns, respectively. Here t_{max} denotes the time of the maximum postsynaptic potential, and the Heaviside step function, defined by $\Theta(x) = 1$ for $x \geq 0$ and $\Theta(x) = 0$ for $x < 0$, ensures zero cost on successful trials. In gradient-based learning, changes in the synaptic efficacies are given by $\Delta\omega_i \propto -dE_\pm/d\omega_i$. Using equation (1) (Results), we found that $dE_\pm/d\omega_i = 0$ for patterns that did not result in erroneous decisions, whereas for erroneous patterns,

$$-\frac{dE_\pm}{d\omega_i} = \pm \sum_{t_i < t_{max}} K(\Delta t_i) \pm \frac{\partial V(t_{max})}{\partial t_{max}} \frac{dt_{max}}{d\omega_i} \qquad (4)$$

The ± sign refers to potentiation and depression following errors on ⊕ and ⊖ patterns, respectively; the sum includes all presynaptic spikes arriving at synapse ω_i before t_{max}, and $\Delta t_i = t_{max} - t_i$ denotes the time difference between the activation of a synapse and the voltage maximum. Because $\partial V(t_{max})/\partial t_{max} = 0$ by definition of t_{max}, only the first term in equation (4) contributes to the synaptic update, yielding the learning rule shown in equation (2). This also holds for ⊖ patterns because of the smooth reset of V following an output spike. We used $\lambda = 10^{-4}/V_0$ for the maximum synaptic change, except for results of the tempotron's performance for different values of τ and N (**Fig. 4**) where we used $\lambda = 3 \times 10^{-3}\, T/(\tau N V_0)$.

Voltage convolution–based learning. In voltage convolution–based learning, changes in synaptic efficacies following error trials were governed by the convolution of the PSP kernel with the postsynaptic potential (equation (3)). Other functional forms that capture the temporal correlation between presynaptic spiking and voltage are possible. The voltage maximum operation was approximated by thresholding υ_i, yielding

$$\Delta\omega_i^\pm = \begin{cases} \pm\lambda & \upsilon_i > \kappa \\ 0 & \upsilon_i \leq \kappa \end{cases}$$

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

for synaptic changes on ⊕ and ⊖ patterns. We used a maximal synaptic change of $\lambda = 8 \times 10^{-5}$ and a plasticity induction threshold of $\kappa = 10^{-3}$. When the system was near its capacity limit, performance could be improved by increasing all weights in which $\upsilon_i \leq \kappa$ by a small amount ($\Delta\omega_i^+ = \varepsilon\lambda$, $\varepsilon = 0.01$) after each erroneous ⊕ pattern trial.

Reinforcement learning. We used a method of reinforcement learning that is based on stochastic synapses (ref. 24). This was implemented by generating, in each trial, independent Gaussian noise ξ_i with mean = 0 and standard deviation = 0.01 and replacing all ω_i in equation (1) with $\omega_i + \xi_i$. Following each error trial, all efficacies were corrected by $\Delta\omega_i = -\lambda\xi_i$.

Momentum term. To accelerate learning, we used a momentum heuristic in all learning rules[15]. In this scheme, synaptic updates $\Delta\omega_i^{\text{current}}$ consist not only of the correction, which is given by the learning rule $\Delta\omega_i$, but also incorporate a fraction of the previous synaptic change $\Delta\omega_i^{\text{previous}}$, such that $\Delta\omega_i^{\text{current}} = \Delta\omega_i + \mu\,\Delta\omega_i^{\text{previous}}$. The parameter μ, which we set to 0.99 throughout the present study, implements a decaying trace of former synaptic changes. As a result, the maximal step size λ was adaptively scaled with a factor ranging from $1/(1-\mu) \approx 100$ if $\Delta\omega_i$ was constant: that is, the learning kept modifying in the same direction, to 1 if the direction of learning oscillated.

Number of contributing synapses. To quantify the number of synapses that contribute to a spike decision at time t_{dec} in response to an input spike pattern, we defined the relative contribution of each synapse, γ_i, as the magnitude of the synaptic efficacy scaled by the PSP kernel, $\gamma_i = |\omega_i|\, K(t_{\text{dec}} - t_i)$. The effective number of synapses contributing to a spike decision was then measured as

$$N_{\text{dec}} = \frac{\left(\sum_{i=1}^{N} \gamma_i\right)^2}{\sum_{i=1}^{N} \gamma_i^2}$$

Note that if M afferents contributed $\gamma_i = \gamma$ and the contribution of the rest was zero, $N_{\text{dec}} = M$ as expected.

Data statistics, parameter choices and initial conditions. Averages over independent realizations were obtained by repeating simulations with identical system parameters but different random number generator seeds (**Figs. 4–6** and **8**). As a result, both input spike patterns and initial synaptic efficacies were different for each realization. When we studied the tempotron near its critical capacity (**Fig. 4**), a small fraction of runs did not converge to zero error within the maximum runtime of our simulations, as expected for any finite N near critical capacity. All data shown complied with the criterion that the probability of misclassifying a pattern after learning was below 10^{-5}.

Additional information on parameter choices, specifications of initial conditions of numerical simulations and Gaussian mixture estimation is presented in **Supplementary Methods** online.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank L. de Hoz, A. Globerson, M. Gutnick, D. Hansel, O. White and Y. Yarom for comments. We acknowledge computational resources provided by the Harvard University Bauer Center for Genomic Research. This work was supported in part by the German Research Foundation, the Minerva Foundation, the European Commission's Improving Human Potential Program, the Israel Science Foundation (Center of Excellence no. 8006/00) and the Defense Research & Development Directorate (MAFAT).

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Published online at http://www.nature.com/natureneuroscience/
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions/

1. Meister, M., Lagnado, L. & Baylor, D.A. Concerted signaling by retinal ganglion cells. *Science* **270**, 1207–1210 (1995).
2. deCharms, R.C. & Merzenich, M.M. Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* **381**, 610–613 (1996).
3. Neuenschwander, S. & Singer, W. Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. *Nature* **379**, 728–732 (1996).
4. Wehr, M. & Laurent, G. Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* **384**, 162–166 (1996).
5. Johansson, R.S. & Birznieks, I. First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat. Neurosci.* **7**, 170–177 (2004).
6. Hopfield, J.J. Pattern recognition computation using action potential timing for stimulus representation. *Nature* **376**, 33–36 (1995).
7. Thorpe, S., Delorme, A. & Van Rullen, R. Spike-based strategies for rapid processing. *Neural Netw.* **14**, 715–725 (2001).
8. Castelo-Branco, M., Goebel, R., Neuenschwander, S. & Singer, W. Neural synchrony correlates with surface segregation rules. *Nature* **405**, 685–689 (2000).
9. Nirenberg, S., Carcieri, S., Jacobs, A. & Latham, P. Retinal ganglion cells act largely as independent encoders. *Nature* **411**, 698–701 (2001).
10. Perkel, D.H. & Bullock, T.H. Neural coding. *Neurosci. Res. Program Bull.* **6**, 221–348 (1968).
11. Singer, W. Time as coding space? *Curr. Opin. Neurobiol.* **9**, 189–194 (1999).
12. Meister, M. & Berry, M. II. The neural code of the retina. *Neuron* **22**, 435–450 (1999).
13. Shadlen, M. & Movshon, J. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**, 67–77 (1999).
14. Minsky, M.L. & Papert, S.A. *Perceptrons* (MIT Press, Cambridge, Massachusetts, 1969).
15. Hertz, J., Krogh, A. & Palmer, R.G. *Introduction to the Theory of Neural Computation* (Westview Press, Boulder, Colorado, 1991).
16. Gardner, E. The space of interactions in neural network models. *J. Phys. A* **21**, 257–270 (1988).
17. Brunel, N., Hakim, V., Isope, P., Nadal, J.P. & Barbour, B. Optimal information storage and the distribution of synaptic weights: perceptron versus purkinje cell. *Neuron* **43**, 745–757 (2004).
18. Vaadia, E. *et al.* Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* **373**, 515–518 (1995).
19. MacLeod, K., Backer, A. & Laurent, G. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* **395**, 693–698 (1998).
20. Brody, C. & Hopfield, J.J. Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron* **37**, 843–852 (2003).
21. Schnitzer, M.J. & Meister, M. Multineuronal firing patterns in the signal from eye to brain. *Neuron* **37**, 499–511 (2003).
22. Kuhn, A., Aertsen, A. & Rotter, S. Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Comput.* **15**, 67–101 (2003).
23. Bressloff, P.B. & Taylor, J.G. Perceptron-like learning in time-summating neural networks. *J. Phys. A* **25**, 4373–4388 (1992).
24. Seung, H.S. Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron* **40**, 1063–1073 (2003).
25. Artola, A., Brocher, S. & Singer, W. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* **347**, 69–72 (1990).
26. Ngezahayo, A., Schachner, M. & Artola, A. Synaptic activity modulates the induction of bidirectional synaptic changes in adult mouse hippocampus. *J. Neurosci.* **20**, 2451–2458 (2000).
27. Sjostrom, P., Turrigiano, G. & Nelson, S. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* **32**, 1149–1164 (2001).
28. Lisman, J. & Spruston, N. Postsynaptic depolarization requirements for LTP and LTD: a critique of spike timing-dependent plasticity. *Nat. Neurosci.* **8**, 839–841 (2005).
29. Cummings, J.A., Mulkey, R.M., Nicoll, R.A. & Malenka, R.C. Ca^{2+} signaling requirements for long-term depression in the hippocampus. *Neuron* **16**, 825–833 (1996).
30. Malenka, R.C. & Nicoll, R.A. Long-term potentiation–a decade of progress? *Science* **285**, 1870–1874 (1999).
31. Kempter, R., Gerstner, W. & van Hemmen, J.L. Hebbian learning and spiking neurons. *Phys. Rev. E* **59**, 4498–4515 (1999).
32. Song, S., Miller, K.D. & Abbott, L.F. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.* **3**, 919–926 (2000).
33. Gütig, R., Aharonov, R., Rotter, S. & Sompolinsky, H. Learning input correlations through non-linear temporally asymmetric Hebbian plasticity. *J. Neurosci.* **23**, 3697–3714 (2003).
34. Guyonneau, R., VanRullen, R. & Thorpe, S.J. Neurons tune to the earliest spikes through STDP. *Neural Comput.* **17**, 859–879 (2005).
35. Legenstein, R., Naeger, C. & Maass, W. What can a neuron learn with spike-timing-dependent plasticity? *Neural Comput.* **17**, 2337–2382 (2005).
36. Worgotter, F. & Porr, B. Temporal sequence learning, prediction, and control: a review of different models and their relation to biological mechanisms. *Neural Comput.* **17**, 245–319 (2005).
37. Foehring, R.C. & Lorenzon, N.M. Neuromodulation, development and synaptic plasticity. *Can. J. Exp. Psychol.* **53**, 45–61 (1999).
38. Seamans, J. & Yang, C. The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Prog. Neurobiol.* **74**, 1–58 (2004).
39. Wespatat, V., Tennigkeit, F. & Singer, W. Phase sensitivity of synaptic modifications in oscillating cells of rat visual cortex. *J. Neurosci.* **24**, 9067–9075 (2004).
40. Lisman, J. A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *Proc. Natl. Acad. Sci. USA* **86**, 9574–9578 (1989).

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

41. Mansuy, I. Calcineurin in memory and bidirectional plasticity. *Biochem. Biophys. Res. Commun.* **311**, 1195–1208 (2003).
42. Liu, L. *et al.* Role of NMDA receptor subtypes in governing the direction of hippocampal synaptic plasticity. *Science* **304**, 1021–1024 (2004).
43. Wang, H.X., Gerkin, R.C., Nauen, D.W. & Bi, G.Q. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.* **8**, 187–193 (2005).
44. Malinow, R. & Malenka, R. AMPA receptor trafficking and synaptic plasticity. *Annu. Rev. Neurosci.* **25**, 103–126 (2002).
45. Zhu, J., Qin, Y., Zhao, M., Van Aelst, L. & Malinow, R. Ras and rap control AMPA receptor trafficking during synaptic plasticity. *Cell* **110**, 443–455 (2002).
46. Tomita, S., Stein, V., Stocker, T., Nicoll, R. & Bredt, D. Bidirectional synaptic plasticity regulated by phosphorylation of stargazin-like tarps. *Neuron* **45**, 269–277 (2005).
47. Lisman, J., Schulman, H. & Cline, H. The molecular basis of CaMKII function in synaptic and behavioural memory. *Nat. Rev. Neurosci.* **3**, 175–190 (2002).
48. Centonze, D., Gubellini, P., Pisani, A., Bernardi, G. & Calabresi, P. Dopamine, acetylcholine and nitric oxide systems interact to induce corticostriatal synaptic plasticity. *Rev. Neurosci.* **14**, 207–216 (2003).
49. Mauk, M. & Buonomano, D. The neural basis of temporal processing. *Annu. Rev. Neurosci.* **27**, 307–340 (2004).
50. Buonomano, D. Decoding temporal information: a model based on short-term synaptic plasticity. *J. Neurosci.* **20**, 1129–1141 (2000).

© 2006 Nature Publishing Group http://www.nature.com/natureneuroscience

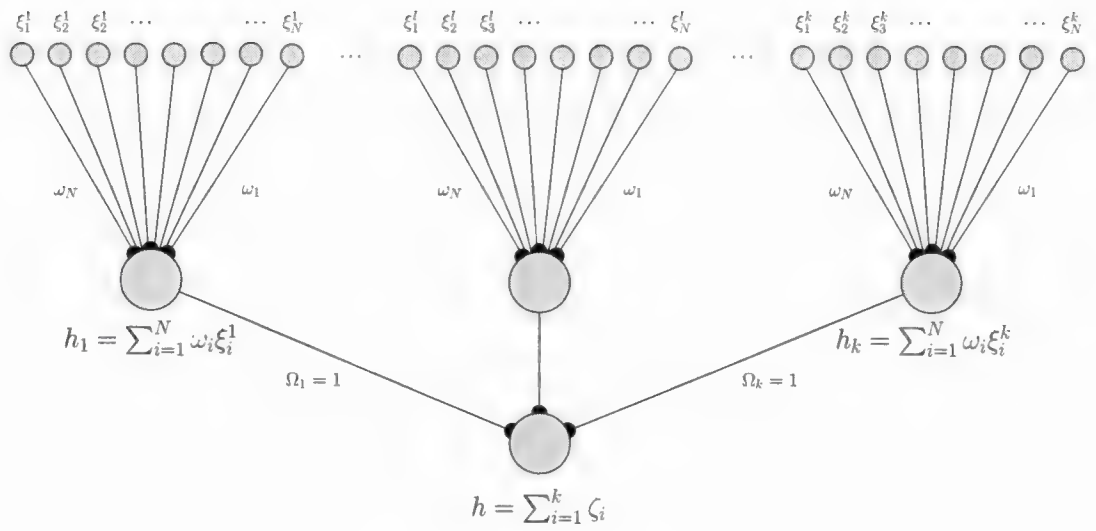
ξ_1^1 ξ_2^1 ξ_2^1 ... ... ξ_N^1
ξ_1^l ξ_2^l ξ_3^l ... ... ξ_N^l
ξ_1^k ξ_2^k ξ_3^k ... ... ξ_N^k
ω_N
ω_1
$h_1 = \sum_{i=1}^N \omega_i \xi_i^1$
$h_k = \sum_{i=1}^N \omega_i \xi_i^k$
$\Omega_1 = 1$
$\Omega_k = 1$
$h = \sum_{i=1}^k \zeta_i$

Supplementary Figure 1, Sompolinsky

Discrete tempotron architecture. A two-layer perceptron which is equivalent to a discrete-time version of the tempotron is shown schematically. The input time interval is segmented into k disjoint 'time slices'. Each integer ξ_i^l denotes the number of spikes that arrive at afferent i during a given time segment l. The variables ξ_i^l, $i = 1, \ldots, N$, are inputs into the l-th perceptron in the next layer. Each of these k perceptrons in this layer has the same set of weights $\omega_1, \ldots, \omega_N$, with which it computes its field h_l and its binary decision ζ_l (Supplementary Note, Eqs. 5,6). The output perceptron sums the ζ_l and computes the final decision σ (Supplementary Note, Eq. 7). Thus an input pattern is classified as $+1$ if at least one ζ_l equals $+1$ and is classified as 0 if they all give 0.

Supplementary Methods

This section gives additional information about the choice of parameters and the initial conditions used in the numerical simulations of Figs. 4-8. We have checked that none of the reported results required fine-tuning of the model parameters.

Number of synapses. All shown simulations except Fig. 4a and Fig. 7c-d were made with $N = 500$. This choice was dictated by computation time and memory constraints. In Fig. 4a, we also show the tempotron's capacity limit for $N = 1000$ and $N = 1500$. In Fig. 7c-d, where each afferent spiked three times, the total number of input spikes was approximately preserved by reducing N to 168. By simulating up to $N = 10,000$ for selected parameters and reduced sample size, we checked that the shown behavior approximates the large N limit.

Time scales. In most of the results, we used the value $\tau = 15\,\mathrm{ms}$ for the PSP decay time constant. This is close to the optimal value (Fig. 4b) when assuming a reasonable pattern duration of $T = 500\,\mathrm{ms}$. For the decay time constant of synaptic currents, we have chosen $\tau_s = \tau/4$, which for $\tau = 15\,\mathrm{ms}$ is compatible with experimental results for fast synapses.

Tempotron learning. In the tempotron learning, we chose the maximal value of synaptic changes to be $\lambda = 10^{-4}/V_0$ (except in Fig. 4). This yields near optimal convergence time with $N = 500$, $T = 500\,\mathrm{ms}$, and $\tau = 15\,\mathrm{ms}$, the parameter valuesq used throughout the work. When τ is far from its optimal, learning times become sensitive to the choice of λ. Empirically we found that $\lambda = 3 \times 10^{-3}\, T/(\tau N V_0)$ results in roughly optimal performance. Hence, this dependence was used in Fig. 4.

Reinforcement learning. We used $\lambda = 10^{-4}$. The standard deviation of synaptic noise (10^{-2}) was chosen as a compromise between minimizing learning time and minimizing the residual error after convergence.

Voltage-convolution based learning. In the voltage-conolution based learning, stable convergence required smaller maximal synaptic changes. We found that the optimal value was roughly $\lambda = 8 \times 10^{-5}$ with $\varepsilon = 0.01$. The plasticity threshold κ, was chosen to be $\kappa = 10^{-3}$. Changing κ slightly degraded the alignment of the mean synaptic change with that of the tempotron (Fig. 8).

Momentum term. Using a momentum term[1] significantly improved convergence times, as has been found in other learning studies. The choice of $\mu = 0.99$ yielded the best results in the present work.

Initial conditions. To avoid encountering pathologies when input spikes arrived in near synchrony (Fig. 4a, gray solid line), a small, hyperpolarizing PSP, peaking at the time of the first input spike, was used to initialize voltages to -0.01. If the subsequent voltage trajectory failed to cross the

resting potential, $t_{\max}$ was set to the time of the last input plus the time to peak of PSPs. In the tempotron learning, initial synaptic efficacies were drawn randomly from a normal distribution with zero mean and standard deviation $\sigma_\omega = 10^{-3}$, except in Fig. 7. In Fig. 7b we found that fast convergence to the solution shown required a larger spread of initial values, $\sigma_\omega = 10^{-2}$. To demonstrate the suppression of synaptic efficacies during the tempotron's acquisition of the multi-neuronal firing pattern task, we also used $\sigma_\omega = 10^{-2}$ in Fig. 7c-d. In voltage-based learning, the mean of the initial distribution of efficacies was chosen to be 10^{-2} to ensure that initial convolved voltages were able to cross the plasticity induction threshold κ. Positive mean of initial efficacies (3×10^{-2}) was also required in the reinforcement algorithm to ensure that the synaptic noise could induce changes in the neuron's response.

Gaussian mixture estimation. Output spike time distributions in Fig. 6b-c were estimated as follows: for each input jitter σ_{in}, output spike times were collected after learning converged from $1,000,000$ randomly drawn templates, resulting in roughly $2,000$ output spikes for each of the 250 $\oplus$ templates used in this task. From these, a histogram was calculated for each template with bin size $\sigma_{\text{in}}/10$ and smoothed with a Gaussian kernel of width $\sigma_{\text{in}}/2$. The distribution of the output spike times from each template was then estimated by the Expectation Maximization algorithm for Gaussian mixture models[2]. The number of components of the mixture model was set to the number of peaks in the smoothed histogram, which was also used to estimate initial values of each components' weight and mean. Initial standard deviations were set to σ_{in}.

1. Hertz, J., Krogh, A. & Palmer, R. G. *Introduction to the Theory of Neural Computation.* (Westview Press, Boulder, 1991).

2. Dempster, N. M., A. P. Laird & Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. B* **39**, 1–38 (1977).

Supplementary Note

Perceptron Learning and Capacity.

The perceptron[1-3] is a node that makes a binary decision by thresholding a weighted sum of its N inputs. For convenience, in this Supplementary Note we denote the decision variables by $\sigma \in \{0,1\}$ and not by $\oplus$ and $\ominus$ as in the paper. The inputs are given by a vector $\vec{x} = (x_1, \ldots, x_N)$, where each component is a single scalar that can be interpreted as the firing rate of an afferent neuron. Thus, the decision of the perceptron is given by

$$\sigma = \Theta\left(\sum_{i=1}^{N} \omega_i x_i - \vartheta\right), \tag{1}$$

where Θ is the Heaviside step function, ω_i denotes the weight of the i-th input component and ϑ is the decision threshold.

A labeled pattern consists of an input vector $\vec{x}$ and a scalar $y \in \{0,1\}$ that denotes the desired output for this input. Given a set of p labeled patterns $\{(\vec{x}^\mu, y^\mu) | \mu = 1, \ldots, p\}$, the perceptron learning rule consists of the following synaptic changes: For each input pattern $\vec{x}^\mu$ for which the perceptron decision σ^μ is not equal to the desired output y^μ, change all weights ω_i $(i = 1, \ldots, N)$ according to

$$\omega_i \mapsto \omega_i + \lambda x_i^\mu (2y^\mu - 1). \tag{2}$$

If a vector of weights that correctly classifies all patterns exists, then this rule converges to some solution vector after a finite number of cycles through all p labeled patterns. It is well-known that this rule converges in a finite number of cycles through all the p labeled patterns to a vector of weights that correctly classify all of them provided that such a vector exists. For large N and random patterns (see below) such a solution exists with probability 1 if and only if $\alpha = p/N \leq \alpha_c$, where $\alpha_c = 2$ (ref.[4]).

The perceptron capacity of $\alpha_c = 2$ holds for patterns that that are randomly labeled $y = 0$ or $y = 1$ with equal probability, but is quite insensitive to the statistics of the inputs x_i. To show this, we consider the case of binary inputs that assume values $\xi_i = 0$ and $\xi_i = 1$ with respective probabilities $(1 + m_{\text{in}})/2$ and $(1 - m_{\text{in}})/2$. Each pattern is labeled as $y = 0$ and $y = 1$ with respective probabilities $(1 + m_{\text{out}})/2$ and $(1 - m_{\text{out}})/2$. Following the derivation of the perceptron capacity in the limit of large N in ref.[4], we obtain for the present case,(cf. Eq. 34-40 in ref.[4])

$$\frac{1}{\alpha_c} = \frac{1}{2}(1 + m_{\text{out}}) \int_x^\infty \mathrm{D}t (t - x)^2 + \frac{1}{2}(1 - m_{\text{out}}) \int_{-x}^\infty \mathrm{D}t (t + x)^2. \tag{3}$$

Here x is an order parameter which is determined through the following saddle-point equation,

$$\frac{1}{2}(1 + m_{\text{out}}) \int_x^\infty \mathrm{D}t (t - x) = \frac{1}{2}(1 - m_{\text{out}}) \int_{-x}^\infty \mathrm{D}t (t + x). \tag{4}$$

This equation determines the value of x and hence of α_c as functions of m_{out} alone, independently of m_{in}. In particular, for $m_{out} = 1/2$, the solution of the above equation is $x = 0$, which upon substitution in Eq. (3) yields $\alpha_c = 2$ for all $0 < m_{in} < 1$. Thus, $\alpha_c = 2$ is a useful benchmark for the capacity of learning random patterns with unbiased labels, such as considered in this paper.

Discrete tempotron

A discrete time version of the tempotron (Supplementary Fig. 1), which is more amenable to statistical mechanical analysis can be devised by segmenting the input time interval $[0, T]$ into k disjoint time slices $[\lambda_l, \lambda_{l+1}]$ with $l = 1, \ldots, k$ and boundaries $\lambda_1 = 0$ and $\lambda_{k+1} = T$. The number of spikes that arrive at afferent i during a given time segment l is denoted by ξ_i^l. For each l the N variables ξ_i^l serve as inputs into a standard perceptron whose output ζ_l is (cf. Eq. 1)

$$\zeta_l = \Theta\left(h_l - \vartheta\right), \tag{5}$$

where the field

$$h_l = \sum_{i=1}^{N} \omega_i \xi_i^l. \tag{6}$$

Note that all k perceptrons in this layer share the same set of weights ω_i. A single second layer perceptron computes the final decision σ by thresholding the sum of all ζ_l,

$$\sigma = \Theta\left(\sum_{l=1}^{k} \zeta_l - \frac{1}{2}\right). \tag{7}$$

Hence, the discrete tempotron gives $\sigma = 1$ if at least one perceptron in the first layer results in $\zeta_l = 1$ and 0 if all first layer perceptrons give 0.

In the discrete tempotron learning is confined to the first layer. By analogy with the continuous tempotron, which aligns synaptic changes to the post-synaptic voltage maximum, the discrete tempotron learning algorithm targets those weights that contribute to the maximal field in the first layer. Specifically, on each error trial the first-layer perceptron l_{max} where

$$l_{max} = \arg\max_l h_l,$$

is subjected to a standard perceptron weight update (cf. Eq. 2) with the desired tempotron output taking the role of the desired single perceptron output.

1. Rosenblatt, F. *Principles of Neurodynamics.* (Spartan, New York, 1962).
2. Minsky, M. L. & Papert, S. A. *Perceptrons* (MIT Press, Cambridge, 1969).
3. Hertz, J., Krogh, A. & Palmer, R. G. *Introduction to the Theory of Neural Computation.* (Westview Press, Boulder, 1991).
4. Gardner, E. The space of interactions in neural network models. *J. Phys. A* **21**, 257–270 (1988).

Time-Warp–Invariant Neuronal Processing

Robert Gütig[1,2]*, Haim Sompolinsky[1,2,3]

1 Racah Institute of Physics, Hebrew University, Jerusalem, Israel, **2** Interdisciplinary Center for Neural Computation, Hebrew University, Jerusalem, Israel, **3** Center for Brain Science, Harvard University, Cambridge, Massachusetts, United States of America

Abstract

Fluctuations in the temporal durations of sensory signals constitute a major source of variability within natural stimulus ensembles. The neuronal mechanisms through which sensory systems can stabilize perception against such fluctuations are largely unknown. An intriguing instantiation of such robustness occurs in human speech perception, which relies critically on temporal acoustic cues that are embedded in signals with highly variable duration. Across different instances of natural speech, auditory cues can undergo temporal warping that ranges from 2-fold compression to 2-fold dilation without significant perceptual impairment. Here, we report that time-warp–invariant neuronal processing can be subserved by the shunting action of synaptic conductances that automatically rescales the effective integration time of postsynaptic neurons. We propose a novel spike-based learning rule for synaptic conductances that adjusts the degree of synaptic shunting to the temporal processing requirements of a given task. Applying this general biophysical mechanism to the example of speech processing, we propose a neuronal network model for time-warp–invariant word discrimination and demonstrate its excellent performance on a standard benchmark speech-recognition task. Our results demonstrate the important functional role of synaptic conductances in spike-based neuronal information processing and learning. The biophysics of temporal integration at neuronal membranes can endow sensory pathways with powerful time-warp–invariant computational capabilities.

Citation: Gütig R, Sompolinsky H (2009) Time-Warp–Invariant Neuronal Processing. PLoS Biol 7(7): e1000141. doi:10.1371/journal.pbio.1000141

Academic Editor: Michael Robert DeWeese, UC Berkeley, United States of America

Received August 13, 2008; **Accepted** May 18, 2009; **Published** July 7, 2009

Copyright: © 2009 Gütig, Sompolinsky. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: RG was funded through fellowships from the Minerva Foundation (Hans-Jensen Fellowship, www.minerva.mpg.de) and the German Science Foundation (GU 605/2-1, www.dfg.de). HS received funding through the Israeli Science Foundation (www.isf.org.il) and MAFAT. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The Hebrew University technology transfer unit (Yissum) has filed the learning rule described in this work and its application to speech processing for patent.

* E-mail: guetig@cc.huji.ac.il

Introduction

Robustness of neuronal information processing to temporal warping of natural stimuli poses a difficult computational challenge to the brain [1–9]. This is particularly true for auditory stimuli, which often carry perceptually relevant information in fine differences between temporal cues [10,11]. For instance in speech, perceptual discriminations between consonants often rely on differences in voice onset times, burst durations, or durations of spectral transitions [12,13]. A striking feature of human performance on such tasks is that it is resilient to a large temporal variability in the absolute timing of these cues. Specifically, changes in speaking rate in ongoing natural speech introduce temporal warping of the acoustic signal on a scale of hundreds of milliseconds, encompassing temporal distortions of acoustic cues that range from 2-fold compression to 2-fold dilation [14,15]. Figure 1 shows examples of time warp in natural speech. The utterance of the word "one" in (A) is compressed by nearly a factor of one-half relative to the utterance shown in (B), causing a concomitant compression in the duration of prominent spectral features, such as the transitions of the peaks in the frequency spectra. Notably, the pattern of temporal warping in speech can vary within a single utterance on a scale of hundreds of milliseconds. For example, the local time warp of the word "eight" in (C) relative to (D), reverses from compression in the initial and final segments to strong dilation of the gap between them. Although it has long been demonstrated that speech perception in humans normalizes durations of temporal cues to the rate of speech [2,16–18], the neural mechanisms underlying this perceptual constancy have remained mysterious.

A general solution of the time-warp problem is to undo stimulus rate variations by comodulating the internal "perceptual" clock of a sensory processing system. This clock should run slowly when the rate of the incoming signal is low and embedded temporal cues are dilated, but accelerate when the rate is fast and the temporal cues are compressed. Here, we propose a neural implementation of this solution, exploiting a basic biophysical property of synaptic inputs, namely, that in addition to charging the postsynaptic neuronal membrane, synaptic conductances modulate its effective time constant. To utilize this mechanism for time-warp robust information processing in the context of a particular perceptual task, synaptic peak conductances at the site of temporal cue integration need to be adjusted to match the range of incoming spike rates. We show that such adjustments can be achieved by a novel conductance-based supervised learning rule. We first demonstrate the computational power of the proposed mechanism by testing our neuron model on a synthetic instantiation of a generic time-warp–invariant neuronal computation, namely, time-warp–invariant classification of random spike latency patterns. We then present a novel neuronal network model for word recognition and show that it yields excellent performance on a benchmark speech-recognition task, comparable to that achieved by highly elaborate, biologically implausible state-of-the-art speech-recognition algorithms.

Author Summary

The brain has a robust ability to process sensory stimuli, even when those stimuli are warped in time. The most prominent example of such perceptual robustness occurs in speech communication. Rates of speech can be highly variable both within and across speakers, yet our perceptions of words remain stable. The neuronal mechanisms that subserve invariance to time warping without compromising our ability to discriminate between fine temporal cues have puzzled neuroscientists for several decades. Here, we describe a cellular process whereby auditory neurons recalibrate, on the fly, their perceptual clocks and allows them effectively to correct for temporal fluctuations in the rate of incoming sensory events. We demonstrate that this basic biophysical mechanism allows simple neural architectures to solve a standard benchmark speech-recognition task with near perfect performance. This proposed mechanism for time-warp–invariant neural processing leads to novel hypotheses about the origin of speech perception pathologies.

Results

Time Rescaling in Neuronal Circuits

Whereas the net current flow into a neuron is determined by the balance between excitatory and inhibitory synaptic inputs, both types of inputs increase the total synaptic conductance, which in turn modulates the effective integration time of the postsynaptic cell [19–21] (an effect known as synaptic shunting). Specifically, when the total synaptic conductance of a neuron is large relative to the resting conductance (leak) and is generated by linear summation of incoming synaptic events, the neuron's effective integration time scales inversely to the rate of inputs spikes. Hence, the shunting action of synaptic conductances can counter variations in afferent spike rates by automatically rescaling the effective integration time of the postsynaptic neuron.

We implement this mechanism in a leaky integrate-and-fire model neuron driven by N exponentially decaying synaptic conductances $g_i(t)=g_i^{\max}\exp(-t/\tau_s)(i=1,\ldots,N)$. Here, $g_i^{\max}$ denotes the peak conductance of the ith synapse in units of sec^{-1}, and τ_s is the synaptic time constant. The total synaptic current, measured at rest, is given by

$$I_{\mathrm{syn}}(t,\beta)=\sum_{i=1}^{N}\sum_{t_i<t}V_i^{\mathrm{rev}}g_i(t-\beta t_i)$$

where V_i^{rev} denotes the reversal potential of the ith synapse relative to resting potential and t_i denote the arrival times of the spikes of the ith afferent. The factor β denotes a global scaling of all incoming spike times; $\beta=1$ is the unwarped inputs. The total synaptic conductance, $G_{\mathrm{syn}}(t,\beta)$, is

$$G_{\mathrm{syn}}(t,\beta)=\sum_{i=1}^{N}\sum_{t_i<t}g_i(t-\beta t_i).$$

For fast synapses, the total synaptic current is essentially a train of pulses, each of which occurs at the time of an incoming spike and delivers a total charge of $g_i\tau_s V_i^{\mathrm{rev}}$. Changing the rate of the incoming spikes will induce a corresponding change in the timing of these pulses but not their charge. Therefore, ignoring the effect of time warp on the time scale of τ_s, which is short relative to the time scale of voltage modulations, the total synaptic current obeys the following time-warp scaling relation, $I_{\mathrm{syn}}(\beta t,\beta)=\beta^{-1}I_{\mathrm{syn}}(t,1)$. A similar scaling relation holds for the total synaptic conductance. The evolution in time of the subthreshold voltage is given by

$$\frac{\mathrm{d}}{\mathrm{d}t}V(t,\beta)=-V(t,\beta)\big(g_{\mathrm{leak}}+G_{\mathrm{syn}}(t,\beta)\big)+I_{\mathrm{syn}}(t,\beta). \quad (1)$$

Thus, V integrates the synaptic current with an effective time constant whose inverse is $1/\tau_{\mathrm{eff}}=g_{\mathrm{leak}}+G_{\mathrm{syn}}(t,\beta)$. If the contribution of G_{syn} is significantly larger than the leak conductance, then $1/\tau_{\mathrm{eff}}$ is rescaled by time-warp similar to G_{syn} and I_{syn}, and, hence, the solution of Equation 1 is approximately time-warp invariant, namely, $V(\beta t,\beta)=V(t,1)$. This result is illustrated in Figure 2, which compares the voltage traces induced by a random spike pattern for $\beta=1$ and $\beta=0.5$.

To perform time-warp–invariant tasks, peak synaptic conductances must be in the range of values appropriate for the statistics of the stimulus ensemble of the given task. To achieve this, we have devised a novel spike-based learning rule for synaptic conductances, the conductance-based tempotron. This model neuron learns to discriminate between two classes of spatiotemporal input spike patterns. The tempotron's classification rule requires it to fire at least one spike in response to each of its target stimuli but to remain silent when driven by a stimulus from the null class. Spike patterns from both classes are iteratively presented to the neuron, and peak synaptic conductances are modified after each error trial by an amount proportional to their contribution to the maximum value of the postsynaptic potential over time (see Materials and Methods). This contribution is sensitive to the time courses of the total conductance and voltage of the postsynaptic neuron. Therefore, the conductance-based tempotron learns to adjust, not only the magnitude of the synaptic inputs, but also its effective integration time to the statistics of the task at hand.

Learning to Classify Time-Warped Latency Patterns

We first quantified the time-warp robustness of the conductance-based tempotron on a synthetic discrimination task. We randomly assigned 1,250 spike pattern templates to target and null classes. The templates consisted of 500 afferents, each firing once at a fixed time chosen randomly from a uniform distribution between 0 and 500 ms. Upon each presentation during training and testing, the templates underwent global temporal warping by a random factor β ranging from compression by $1/\beta_{\max}$ to dilation by $\beta_{\max}$ (see Materials and Methods). Consistent with the psychophysical range, $\beta_{\max}$ was varied between 1 and 2.5. Remarkably, with physiologically plausible parameters, the error frequency remained almost zero up to $\beta_{\max}\approx 2$ (Figure 3A, blue curve). Importantly, the performance of the conductance-based tempotron showed little change when the temporal warping applied to the spike templates was dynamic (see Materials and Methods) (Figure 3A). The time-warp robustness of the neural classification depends on the resting membrane time constant τ_m and the synaptic time constant τ_s. Increases in τ_m or decreases in τ_s both enhance the dominance of shunting in governing the cell's effective time constant. As a result, the performance for $\beta_{\max}=2.5$ improved with increasing τ_m (Figure 3B, left) and decreasing τ_s (Figure 3B, right). The time-warp robustness of the conductance-based tempotron was also reflected in the shape of its subthreshold voltage traces (Figure 3C, top row) and generalized to novel spike templates with the same input statistics that were not used during training (Figure 3C, second row).

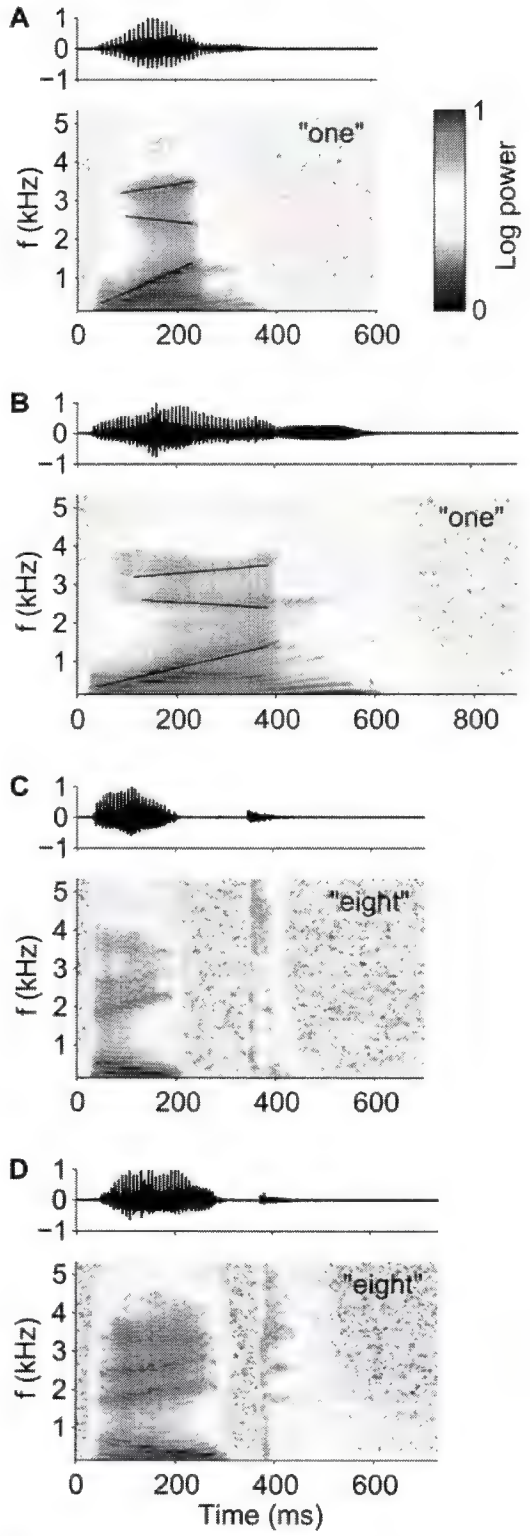

Figure 1. Time warp in natural speech. Sound pressure waveforms (upper panels, arbitrary units) and spectrograms (lower panels, color-code scaled between the minimum and maximum log power) of speech samples from the TI46 Word corpus [24], spoken by different male speakers. (A and B) Utterances of the word "one." Thin black lines highlight the transients of the second, third, and fourth (bottom to top) spectral peaks (formants). The lines in (A) are compressed relative to (B) by a common factor of 0.53. (C and D) Utterances of the word "eight."
doi:10.1371/journal.pbio.1000141.g001

Synaptic conductances were crucial in generating the neuron's robustness to temporal warping. Athough an analogous neuron model with a fixed integration time, the current-based tempotron [22] (see Materials and Methods) also performed the task perfectly in the absence of time-warp ($\beta_{max} = 1$); its error frequency was sensitive even to modest temporal warping and deteriorated further when the applied time warp was dynamic (Figure 3A, red curve). Similarly, the voltage traces of this current-based neuron showed strong dependence on the degree of temporal warping applied to an input spike train (Figure 3C, bottom trace pair). Finally, the error frequency of the current-based neuron at $\beta_{max} = 2.5$ showed only negligible improvement upon varying the values of the membrane and synaptic time constants (Figure 3B), highlighting the limited capabilities of fixed neural kinetics to subserve time-warp–invariant spike-pattern classification.

Note that in the present classification task, the degree of time-warp robustness depends also on the learning load, i.e., number of

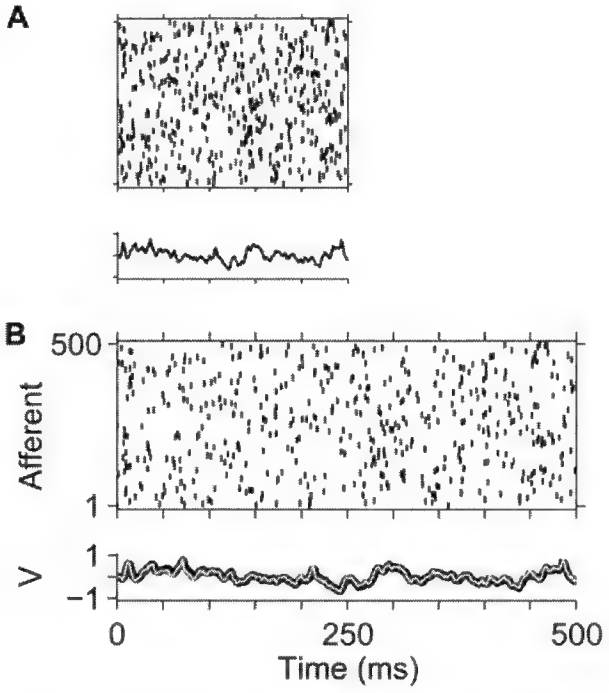

Figure 2. Time-warp–invariant voltage traces. Spike rasters show a random spike pattern across $N = 500$ afferents ($N_{ex} = 250$ excitatory and $N_{in} = 250$ inhibitory), each of which fires a single action potential at a random time chosen uniformly between 0 and 500 ms. Whereas the original spike pattern ($\beta = 1$) is shown in (B), the pattern displayed in (A) is compressed by a factor of $\beta = 0.5$. In each panel, the lower trace depicts the voltage $V(t,\beta)$ induced by the spike patterns in our model neuron with balanced uniform synaptic peak conductances that resulted in a zero mean synaptic current at rest set to $g_{ex}^{max} = 6/(N_{ex}\tau_s)$ for excitatory synapses and $g_{in}^{max} = 5g_{ex}^{max}$ for inhibitory synapses. These values result in a mean total synaptic conductance of $\overline{G}_{syn} \approx 7g_{leak}$. In (B), the voltage trace $V(t,1)$ (thin grey line) is superimposed on the rescaled voltage trace $V(\beta t,\beta)$ (thick black line) from (A).
doi:10.1371/journal.pbio.1000141.g002

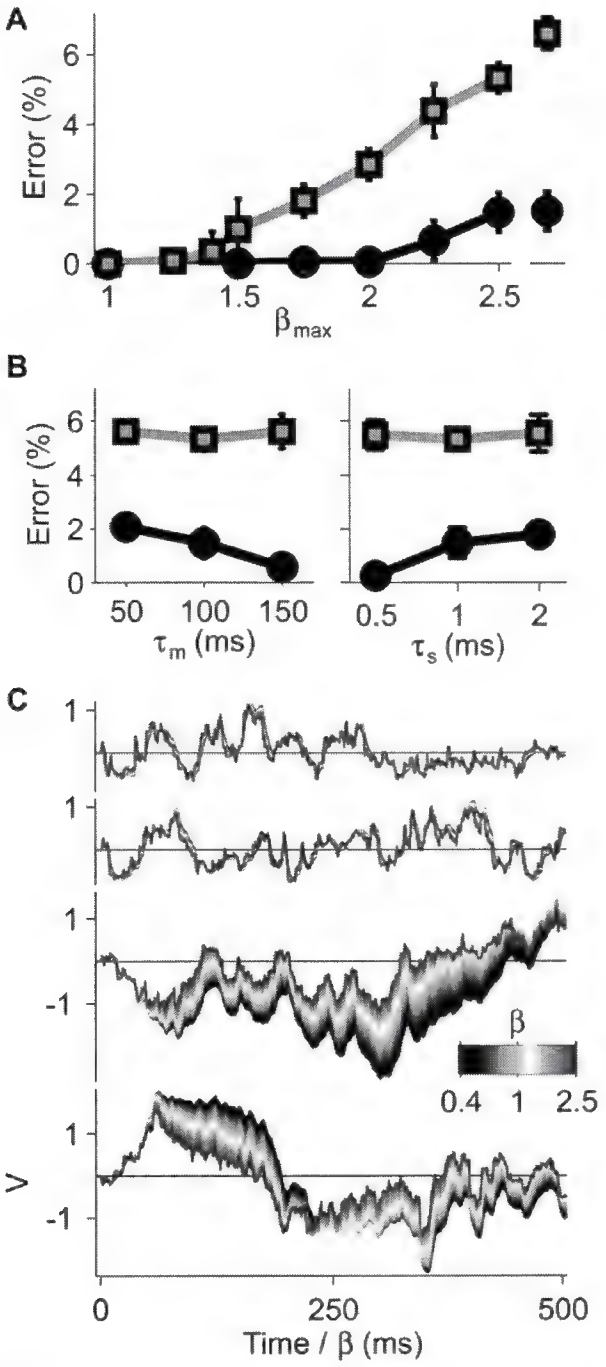

Figure 3. Classification of time-warped random latency patterns. (A) Error probabilities versus the scale of global time-warp β_{max} for the conductance-based (blue) and the current-based (red) neurons. Errors were averaged over 20 realizations, error bars depict ±1 standard deviation (s.d.). Isolated points on the right were obtained under dynamic time warp with $\beta_{max}=2.5$ (see Materials and Methods). (B) Dependence of the error frequency at $\beta_{max}=2.5$ on the resting membrane time constant τ_m (left) and the synaptic time constant τ_s (right). Colors and statistics as in (A). (C) Voltage traces of a conductance-based (top and second rows) and a current-based neuron (third and bottom rows). Each trace was computed under global time warp with a temporal scaling factor β (see Materials and Methods) (color bar) and plotted versus a common rescaled time axis. For each neuron model, the upper traces were elicited by a target and the lower traces by an untrained spike template.
doi:10.1371/journal.pbio.1000141.g003

patterns that have to classified by a neuron (unpublished data). A given degree of time warp translates into a finite range of distortions of the intracellular voltage traces. If these distortions remain smaller than the margins separating the neuronal firing threshold and the intracellular peak voltages, a neuron's classification will be time-warp invariant. Since the maximal possible margins increase with decreasing learning load, time-warp invariance can be traded for storage capacity. This tradeoff is governed by the susceptibility of the voltage traces to time warp. If the susceptibility is high, as in the current-based tempotron, robustness to time warp comes at the expense of a substantial reduction in storage capacity. If it is low, as in the conductance-based tempotron, time-warp invariance can be achieved even when operating close to the neuron's maximal storage capacity for unwarped patterns.

Adaptive Plasticity Window

In the conductance-based tempotron, synaptic conductances controlled, not only the effective integration time of the neuron, but also the temporal selectivity of the synaptic update during learning. The tempotron learning rule modifies only the efficacies of the synapses that were activated in a temporal window prior to the peak in the postsynaptic voltage trace. However, the width of this temporal plasticity window is not fixed but depends on the effective integration time of the postsynaptic neuron at the time of each synaptic update trial, which in turn varies with the input firing rate at each trial and the strength of the peak synaptic conductances at this stage of learning (Figure 4). During epochs of high conductance (warm colors), only synapses that fired shortly before the voltage maximum were appreciably modified. In contrast, when the membrane conductance was low (cool colors), the plasticity window was broad. The ability of the plasticity window to adjust to the effective time constant of the postsynaptic voltage is crucial for the success of the learning. As is evident from Figure 4, the membrane's effective time constant varies consider-

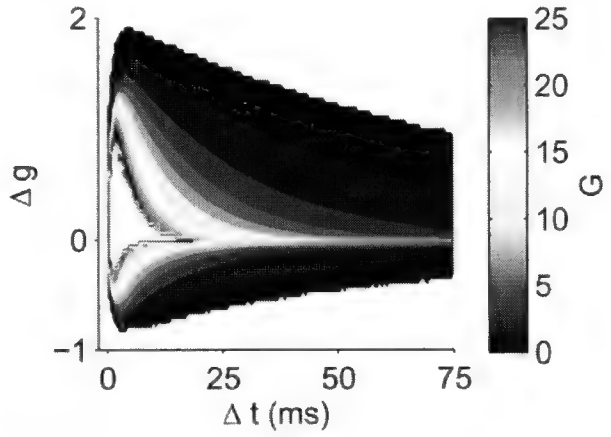

Figure 4. Adaptive learning kernel. Change in synaptic peak conductance Δg versus the time difference Δt between synaptic firing and the voltage maximum, as a function of the mean total synaptic conductance G during this interval (color bar). Data were collected during the initial 100 cycles of learning with $\beta_{max}=2.5$ and averaged over 100 realizations.
doi:10.1371/journal.pbio.1000141.g004

ably during the learning epochs; hence, a plasticity rule that does not take this into account fails to credit appropriately the different synapses.

Task Dependence of Learned Synaptic Conductance

The evolution of synaptic peak conductances during learning was driven by task requirements. When we replaced the temporal warping of the spike templates by random Gaussian jitter [22] (see Materials and Methods), conductance-based tempotrons that had acquired high synaptic peak conductances during initial training on the time-warp task readjusted their synaptic peak conductances to low values (Figure 5, inset). The concomitant increase in their effective integration time constants from roughly 10 ms to 50 ms improved the neurons' ability to average out the temporal spike jitter and substantially enhanced their task performance (Figure 5).

Neuronal Model of Word Recognition

To address time-warp–invariant speech processing, we studied a neuronal module that learns to perform word-recognition tasks. Our model consists of two auditory processing stages. The first stage (Figure 6) consists of an afferent population of neurons that convert incoming acoustic signals into spike patterns by encoding the occurrences of elementary spectrotemporal events. This layer forms a 2-dimensional tonotopy-intensity auditory map. Each of

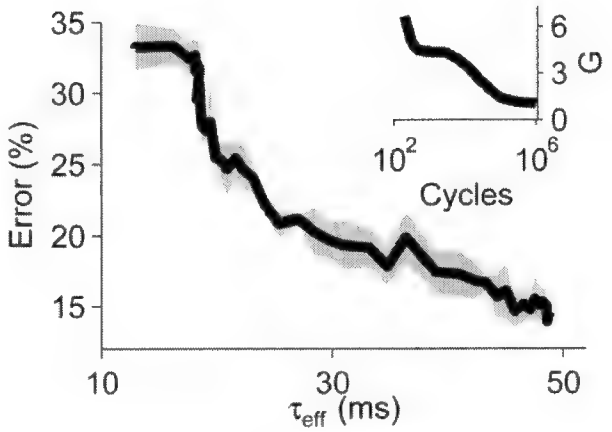

Figure 5. Task dependence of the learned total synaptic conductance. Error frequency of the conductance-based tempotron versus its effective integration time τ_{eff}. After switching from time-warp to Gaussian spike jitter, τ_{eff} increased as the mean time-averaged total synaptic conductance G decreased with learning time (inset).
doi:10.1371/journal.pbio.1000141.g005

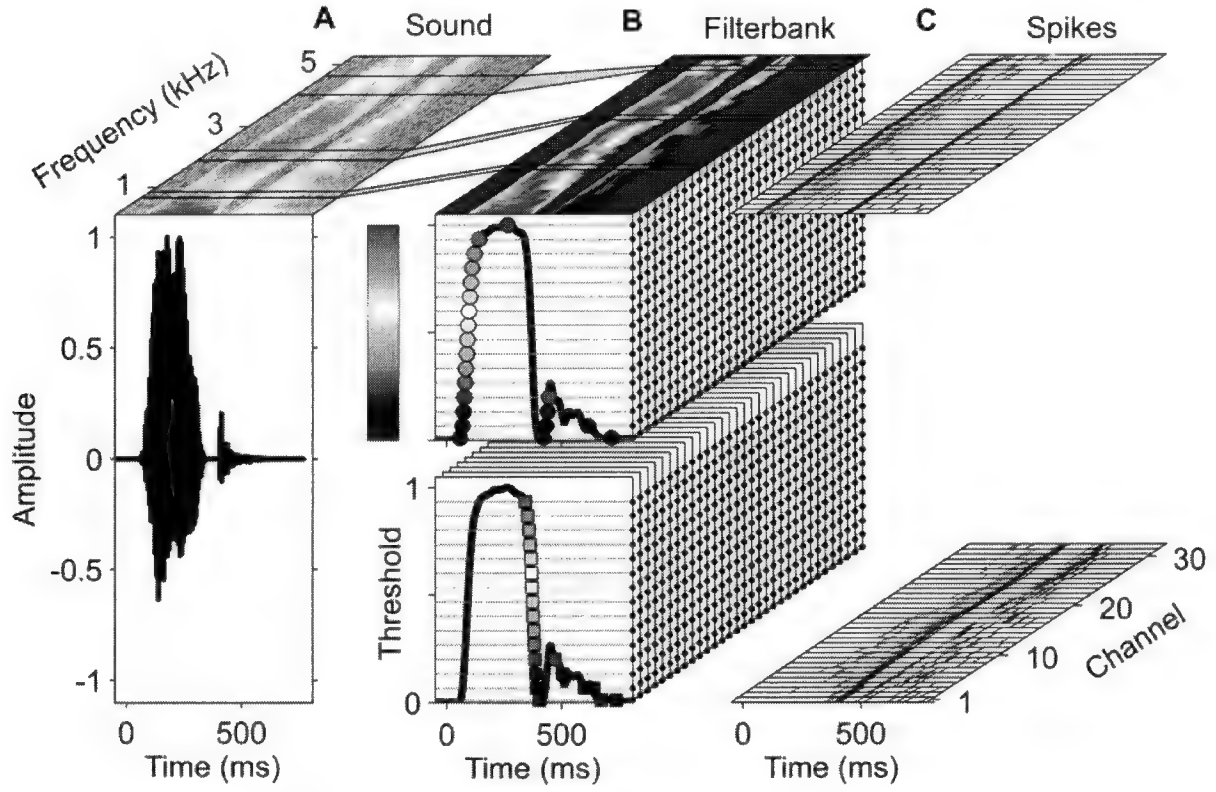

Figure 6. Auditory front end. (A and B) Incoming sound signal (bottom) and its spectrogram in linear scale (top) as in Figure 1D (A). Based on the spectrogram, the log signal power in 32 frequency channels (Mel scale, see Materials and Methods) is computed and normalized to unit peak amplitude in each channel ([B], top, colorbar). Black lines delineate filterbank channels 10, 20, and 30 and their respective support in the spectrogram (connected through grey areas). In each channel, spikes in 31 afferents (small black circles) are generated by 16 onset (upper block) and 15 offset (lower block) thresholds. For the signal in channel 1 (shown twice as thick black curves on the front sides of the upper and lower blocks), resulting spikes are marked by circles (onset) and squares (offset) with colors indicating respective threshold levels (colorbar). (C) Spikes (onset, top, and offset, bottom) from all 992 afferents plotted as a function of time (*x*-axis) and corresponding frequency channel (*y*-axis). The color of each spike (short thin lines) indicates the threshold level (as used for circles and squares in [B]) of the eliciting unit.
doi:10.1371/journal.pbio.1000141.g006

its afferents generates spikes by performing an onset or offset threshold operation on the power of the acoustic signal in a given frequency band. Whereas an onset afferent elicits a spike whenever the log signal power crosses its threshold level from below, offset afferents encode the occurrences of downward crossings (see Materials and Methods) (cf. also [6,23]). Different on and off neurons coding for the same frequency band differ in their threshold value, reflecting a systematic variation in their intensity tuning. The second, downstream, layer consists of neurons with plastic synaptic peak conductances that are governed by the conductance-based tempotron plasticity rule. These neurons are trained to perform word discrimination tasks. We tested this model on a digit-recognition benchmark task with the TI46 database [24]. We trained each of the 20 conductance-based tempotrons of the second layer to perform a distinct gender-specific binary classification, requiring it to fire in response to utterances of one digit and speaker gender, and to remain quiescent for all other stimuli. After training, the majority of these digit detector neurons (70%) achieved perfect classification of the test set, and the remaining ones performed their task with a low error (Table 1). Based on the spiking activity of this small population of digit detector neurons, a full digit classifier (see Materials and Methods) that weighted spikes according to each detector's individual performance, achieved an overall word error rate of 0.0017. This performance matches the error rates of state-of-the-art artificial speech-recognition systems such as the Hidden Markov model–based Sphinx-4 and HTK, which yield error rates of 0.0017 [25] and 0.0012 [26], respectively, on the same benchmark.

Table 1. Test set error fractions of individual detector neurons.

Digit	Male	Female
0	0.0	0.0
1	0.0	0.0
2	0.0008	0.0017
3	0.0	0.0
4	0.0	0.0
5	0.0029	0.0062
6	0.0	0.0
7	0.0004	0.0008
8	0.0	0.0
9	0.0	0.0

doi:10.1371/journal.pbio.1000141.t001

Learned Spectrotemporal Target Features

To reveal qualitatively some of the mechanisms used by our digit detector neurons to selectively detect their target word, we compared the learned synaptic distributions (Figure 7A) of two digit detector neurons ("one" and "four") to the average spectrograms of each neuron's target stimuli aligned to the times of its output spikes (Figure 7B; see Materials and Methods). The spectrotemporal features that preceeded the output spikes (time zero, grey vertical lines) corresponded to the frequency-specific onset and offset selectivity of the excitatory afferents (Figure 7A, warm colors). These examples demonstrate how the patterned excitatory and inhibitory inputs from both onset and offset neurons are tuned to features of the speech signal. For instance, a prominent feature in the averaged spectrogram of the word "one" (male speakers) was the increase in onset time of the power in the low-frequency channels with the frequency of the channel (Figure 7B, left, channels 1–16). This gradual onset was encoded by a diagonal band of excitatory onset afferents whose thresholds decreased with increasing frequency (Figure 7A, left). By compensating for the temporal lag between the different lower-frequency channels, this arrangement ensured a strong excitatory drive when a target stimulus was presented to the neuron. The spectrotemporal feature learned by the word "four" (male speakers) detector neuron combined decreasing power in the low-frequency range with rising power in the mid-frequency range (Figure 7B, right). This feature was encoded by synaptic efficacies through a combination of excitatory offset afferents in the low-frequency range (Figure 7A, right, channels 1–11) and excitatory onset afferents in the mid-frequency range (channels 12–19). Excitatory synaptic populations were complemented by inhibitory inputs (Figure 7A, blue patches) that prevented spiking in response to null stimuli and also increased the total synaptic conductance. The substantial differences between the mean spike-triggered voltage traces for target stimuli (Figure 7C, blue) and the mean maximum-triggered voltage traces for null stimuli (red) underline the high target word selectivity of the learned synaptic distributions as well as the relatively short temporal extend of the learned target features.

In the examples shown, the average position of the neural decision relative to the target stimuli varied from early to late (Figure 7B, left vs. right). This important degree of freedom stems from the fact that the tempotron decision rule does not constrain the time of the neural decision. As a result, the learning process in each neuron can select the spectrotemporal target features from any time window within the word. The selection of the target feature by the learning takes into account both the requirement of triggering output spikes in response to target stimuli as well as the demand to remain silent during null stimuli. Thus, for each target neuron, the selected features reflect the statistics of both the target and the null stimuli.

Generalization Abilities of Word Detector Neurons

We have performed several tests designed to assess the ability of the model word detector neurons to perform well on new input sets, different in statistics from the trained database. First, we assessed the ability of the neurons to generalize to unfamiliar speakers and dialects. After training the model with the TI46 database, as described above, we measured its digit-recognition performance on utterances from another database, the TIDIGITS database [27], which includes speech samples from a variety of English dialects (see Materials and Methods). This test has been done without any retraining of the network synapses. The resulting word error rate of 0.0949 compares favorably to the performance of the HTK system, which resulted in an error rate of 0.2156 when subjected to the same generalization test (see Materials and Methods). Across all dialects, our model performed perfectly for roughly one-quarter of all speakers and with at most one error for half of them. Within the best dialect group, an error of at most one word was achieved for as many as 80% of the speakers (Table S1). These results underline the ability of our neuronal word-recognition model to generalize to unfamiliar speakers across a wide range of different unfamiliar dialects.

An interesting question is whether our model neurons are able to generalize their performance to novel time-warped versions of the trained inputs. To address this question, we have tested their performance on randomly generated time-warped versions of the input spikes corresponding to the trained word utterances, without retraining. As shown in Figure 8, the neurons exhibited considerable time-warp–robust performance on the digit-recogni-

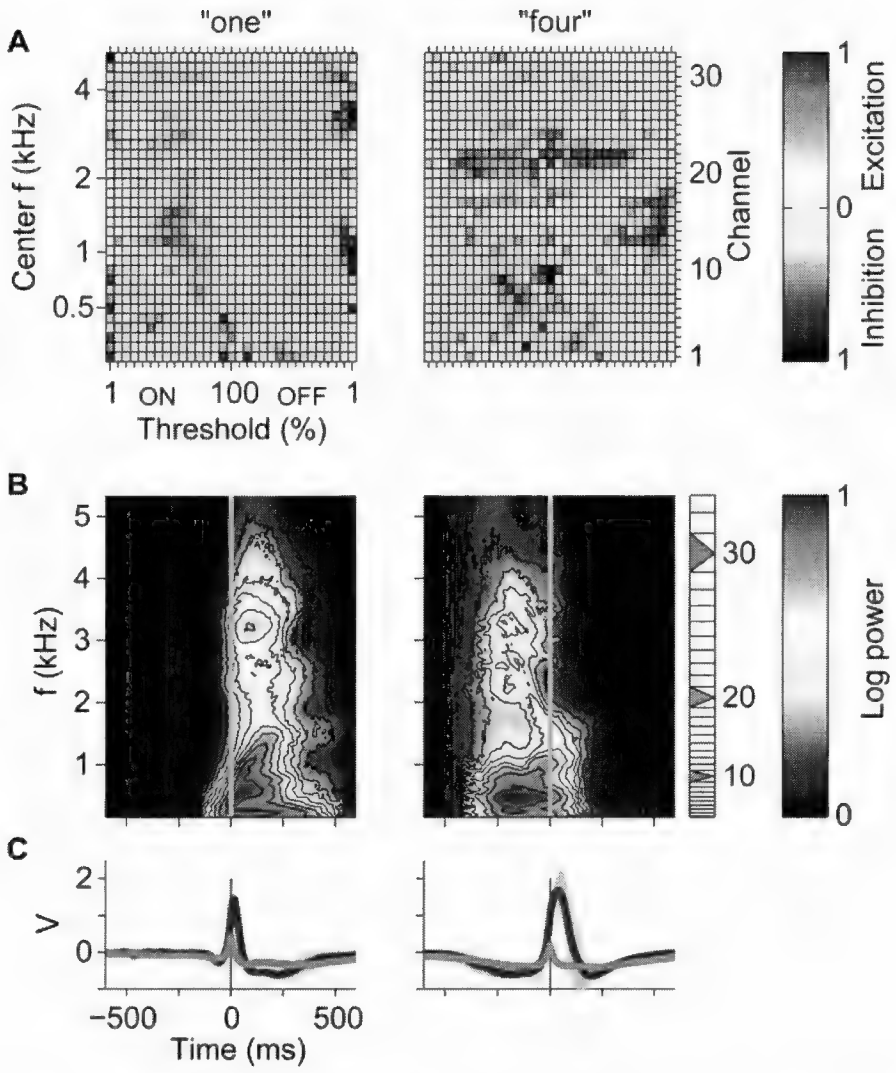

Figure 7. Speech-recognition task. (A) Learned synaptic peak conductances. Each pixel corresponds to one synapse characterized by its frequency channel (right *y*-axis) and its onset (ON) or offset (OFF) afferent power threshold level (*x*-axis, in percent of maximum signal powers [see Materials and Methods]). Learned peak conductances were color coded with excitatory (warm colors) and inhibitory conductances (cool colors) separately normalized to their respective maximal values (color bar). The left *y*-axis shows the logarithmically spaced center frequencies (Mel scale) of the frequency channels. (B) Spike-triggered target stimuli (color-code scaled between the minimum and maximum mean log power). (C) Mean voltage traces for target (blue, light blue ±1 s.d.; spike triggered) and null stimuli (red; maximum triggered).
doi:10.1371/journal.pbio.1000141.g007

tion task. For instance, the errors for the "one" (Figure 8A, black line) and "four" (blue line) detector neurons (cf. Figure 7) were insensitive to a 2-fold time warp of the input spike trains. The "seven" detector neuron (male, red line) showed higher sensitivity to such warping; nevertheless, its error rate remained low. Consistent with the proposed role of synaptic conductances, the degree of time-warp robustness was correlated with the total synaptic conductance, here quantified through the mean effective integration time τ_{eff} (Figure 8B). Additionally, the mean voltage traces induced by the target stimuli (Figure 8C, lower traces) showed a substantially smaller sensitivity to temporal warping than their current-based analogs (see Materials and Methods) (Figure 8C, upper traces).

We also found that our model word detector neurons are robust to the introduction of spike failures in their input patterns. For each neuron, we have measured its performance on inputs which were corrupted by randomly deleting a fraction of the incoming spikes, again without retraining. For the majority of neurons, the error percentage increased by less than 0.01% for each percent increase in spike failures (Figure 9). This high robustness reflects the fact that each classification is based on integrating information from many presynaptic sources.

Discussion

Automatic Rescaling of Effective Integration Time by Synaptic Conductances

The proposed conductance-based time-rescaling mechanism is based on the biophysical property of neurons that their effective integration time is shaped by synaptic conductances and therefore

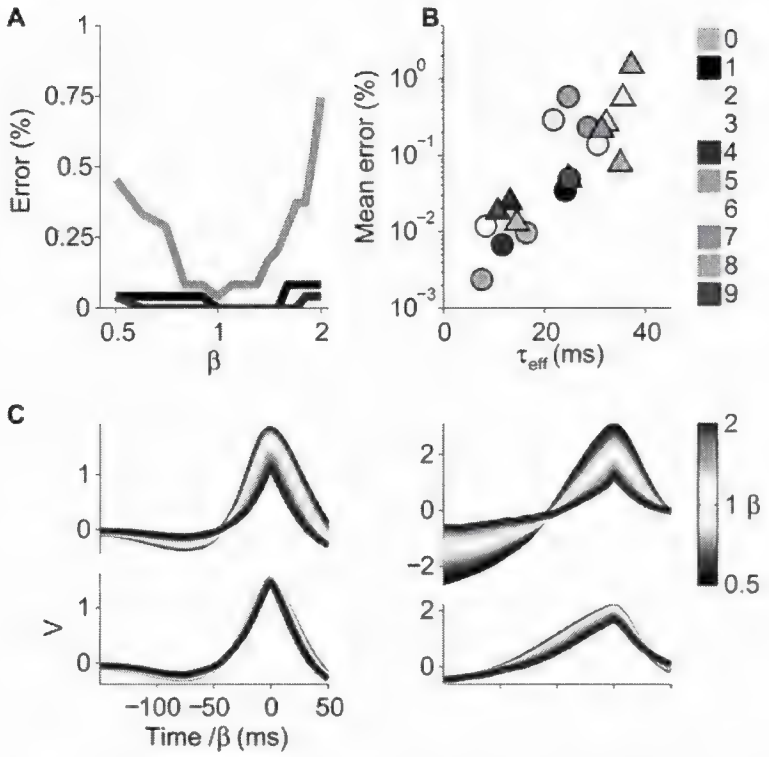

Figure 8. Time-warp robustness. (A) Error versus time-warp factor β. (B) Mean errors over the range of β shown in (A) (digit color code; triangles: female speakers, circles: male speakers) versus the mean effective time constant τ_{eff} calculated for $\beta = 1$ by averaging the total synaptic conductance over 100-ms time windows prior to either the output spikes (target stimuli) or the voltage maxima (null stimuli). (C) Mean voltage traces for time-warped target patterns for the neurons shown in Figure 7. Bottom row: conductance-based neurons, upper row: current-based neurons (see Materials and Methods).
doi:10.1371/journal.pbio.1000141.g008

can be modulated by the firing rate of its afferents. To utilize these modulations for time-warp–invariant processing, a central requirement is a large evoked total synaptic conductance that dominates the effective integration time constant of the postsynaptic cell through shunting. In our speech-processing model, large synaptic conductances with a median value of a 3-fold leak conductance across all digit detector neurons (cf. Figure 8B) result from a combination of excitatory and inhibitory inputs. This is consistent with high total synaptic conductances, comprising excitation and inhibition, that have been observed in several regions of cortex [28] including auditory [29,30], visual [31,32], and also prefrontal [33,34] (but see ref. [35]). Our model predicts that in cortical sensory areas, the time-rescaled intracellular voltage traces (cf. Figure 3C), and consequently, also the rescaled spiking responses of neurons that operate in the proposed fashion, remain invariant under temporal warping of the neurons' input spike patterns. These predictions can be tested by intra- and extracellular recordings of neuronal responses to temporally warped sensory stimuli.

A large total synaptic conductance is associated with a substantial reduction in a neuron's effective integration time relative to its resting value. Therefore, the resting membrane time constant of a neuron that implements the automatic time-rescaling mechanism must substantially exceed the temporal resolution that is required by a given processing task. Because the word-recognition benchmark task used here comprises whole-word stimuli that favored effective time constants on the order of several tens of milliseconds, we used a resting membrane time constant of $\tau_m = 100$ ms. Whereas values of this order have been reported in hippocampus [36] and cerebellum [21,37], it exceeds current estimates for neocortical neurons, which range between 10 and 30 ms [35,38,39]. Note, however, that the correspondence of our passive membrane model and the experimental values that typically include contributions from various voltage-dependent conductances is not straightforward. Our model predicts that neurons specialized for time-warp–invariant processing at the whole-word level have relatively long resting membrane time constants. It is likely that the auditory system solves the problem of time-warp–invariant processing of the sound signal primarily at the level of shorter speech segments such as phonemes. This is supported by evidence that primary auditory cortex has a special role in speech processing at a resolution of milliseconds to tens of milliseconds [11–13]. Our mechanism would enable time-warp–invariant processing of phonetic segments with resting membrane time constants in the range of tens of milliseconds, and much shorter effective integration times.

The proposed neuronal time-rescaling mechanism assumes linear summation of synaptic conductances. This assumption is challenged by the presence of voltage-dependent conductances in neuronal membranes. Since the potential implications for our model depend on the specific nonlinearity induced by a cell-type–specific composition of different ionic channels, it is hard to evaluate the overall effect on our model in general terms. Nevertheless, because of its immanence, we expect the conductance-based time-rescaling mechanism to cope gracefully with moderate levels of nonlinearity. As an example, we tested its

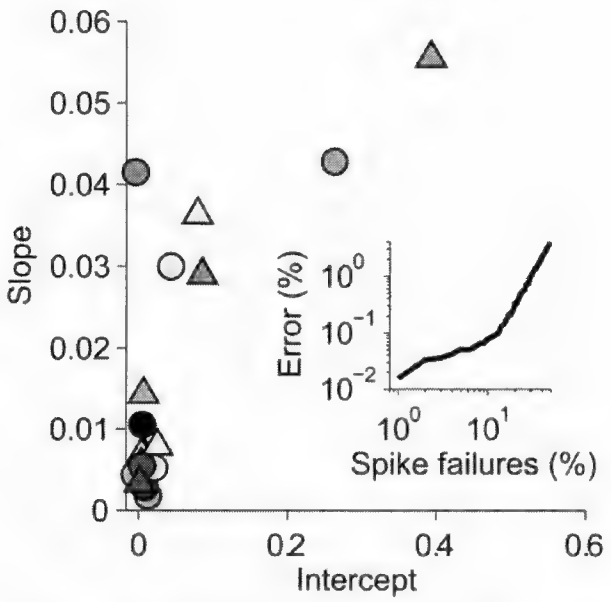

Figure 9. Robustness to spike failures. The error fraction of each digit detector neuron was measured as a function of the spike failure probability over the range from 0% to 10% and fitted by linear regression. For each neuron, the resulting slope (median 0.0069) is plotted versus the intercept (median 0.0061) with symbols and colors as in Figure 8B. The median R^2 of the linear regression fits was 0.94. The inset shows the median error fraction of the population as a function of the spike failure probability in the range of 1% to 50% with the robust regime braking down at approximately 20%.
doi:10.1371/journal.pbio.1000141.g009

behavior in the presence of an h-like conductance (see Materials and Methods) that opposes conductance changes induced by depolarizing excitatory synaptic inputs and is active at the resting potential. As expected, we found that physiological levels of h-conductances resulted in only moderate impairment of the automatic time-rescaling mechanism (Figure S1).

For the sake of simplicity as well as numerical efficiency, we have assumed symmetric roles of excitation and inhibition in our model architecture. We have checked that this assumption is not crucial for the operation of the automatic time-rescaling mechanism and the learning of time-warped random latency patterns. Specifically, we have implemented the random latency classification task for a control architecture in which all synapses were confined to be excitatory except a single global inhibitory input that, mimicking a global inhibitory network, received a separate copy of all incoming spikes. In this architecture, all spike patterns have to be encoded by the excitatory synaptic population, and the role of inhibition is reduced to a global signal that has equal strength for all input patterns. Due to the limitations of this architecture, this model showed some reduction of storage capacity relative to the symmetric case, but the automatic time-rescaling mechanism remained intact. For a time-warp scale of $\beta_{max} = 2.5$ (cf. Figure 3), the global inhibition model roughly matched the performance of the symmetric model when the learning load was lowered to 1.5 spike patterns per synapse, with an error fraction of 0.18%.

Supervised Learning of Synaptic Conductances

To utilize synaptic conductances as efficient controls of the neuron's clock, the peak synaptic conductances must be plastic so that they adjust to the range of integration times relevant for a given perceptual task. This was achieved in our model by our novel supervised spike-based learning rule. This plasticity posits that the temporal window during which pre- and postsynaptic activity interact continuously adapts to the effective integration time of the postsynaptic cell (Figure 4). The polarity of synaptic changes is determined by a supervisory signal that we hypothesize to be realized through neuromodulatory control [22]. Because present experimental measurements of spike-timing–dependent synaptic plasticity rules have assumed an unsupervised setting, i.e., have not controlled for neuromodulatory signals (but see [40]), existing results do not directly apply to our model. Nevertheless, recent data have revealed complex interactions between the statistics of pre- and postsynaptic spiking activity and the expression of synaptic changes [41–44]. Our model offers a novel computational rationale for such interactions, predicting that for fixed supervisory signaling, the temporal window of plasticity shrinks with growing levels of postsynaptic shunting. One challenge for the biological implementation of the tempotron learning rule is the need to compute the time of the maximum of the postsynaptic voltage. We have previously shown for a current-based neuron model that this temporally global operation can be approximated by temporally local computations that are based on the postsynaptic voltage traces following input spikes [22]. We have extended this approach to plastic synaptic conductances and checked that the resulting biologically plausible implementation of conductance-based tempotron learning can readily subserve time-warp–invariant classification of spike patterns. Specifically, in this implementation, the induction of synaptic plasticity is controled by the correlation of the postsynaptic voltage and a synaptic learning kernel (see Materials and Methods) whose temporal extend is controlled by the average conductance throughout a given error trial. A synaptic peak conductance is changed by a uniform amount whenever this correlation exceeds a fixed plasticity induction threshold. When tested on the time-warped latency patterns with $\beta_{max} = 2.5$ (cf. Figure 3), the correlation-based tempotron roughly matched the voltage maximum–based version at a reduced learning load of 1.5 patterns per synapse with an error fractions of 0.35%.

Time-Warp Invariance Is Task Dependent

In our model, dynamic time-warp–invariant capabilities become avaliable through a conductance-based learning rule that tunes the shunting action of synaptic conductances. This learning rule enables neurons to adjust the degree of synaptic shunting to the requirements of a given processing task. As a result, our model can naturally encompass a continuum of functional specializations ranging from neurons that are sensitive to absolute stimulus durations by employing low total synaptic conductances, to time-warp–invariant feature detectors that operate in a high-conductance regime. In the context of auditory processing, such a functional segregation into neurons with slower and faster effective integration times is reminiscent of reports suggesting that rapid temporal processing in time frames of tens of milliseconds is localized in left lateralized language areas, whereas processing of slower temporal features is attributed to right hemispheric areas [45–47]. Although anatomical and morphological asymmetries between left and right human auditory cortices are well documented [48], it remains to be seen whether these differences form the physiological substrate for a left lateralized implementation of the proposed time-rescaling mechanism. Consistent with this picture, the general tradeoff between high temporal resolution and robustness to temporal jitter that is predicted by our model (Figure 5) parallels reports of the vulnerability of the lateralizion of

language processing with respect to background acoustic noise [49] as well as to abnormal timing of auditory brainstem responses [50].

Neuronal Circuitry for Time-Warp–Invariant Feature Detection

The architecture of our speech-processing model encompasses two auditory processing stages. The first stage transforms acoustic signals into spatiotemporal patterns of spikes. To engage the proposed automatic time-rescaling mechanism, the population rate of spikes elicited in this afferent layer must track variations in the rate of incoming speech. Such behavior emerges naturally in a sparse coding scheme in which each neuron responds transiently to the occurrences of a specific acoustic event within the auditory input. As a result, variations in the rate of acoustic events are directly translated into concomitant variations in the population rate of elicited spikes. In our model, the elementary acoustic events correspond to onset and offset threshold crossings of signal power within specific frequency channels. Such frequency-tuned onset and offset responses featuring a wide range of dynamic thresholds have been observed in the inferior colliculus (IC) of the auditory midbrain [51]. This nucleus is the site of convergence of projections from the majority of lower auditory nuclei and is often referred to as the interface between the lower brain stem auditory pathways and the auditory cortex. Correspondingly, we hypothesize that the layer of time-warp–invariant feature detector neurons in our model implements neurons located downstream of the IC, most probably in primary auditory cortex. Current studies on the functional role of the auditory periphery in speech perception and its pathologies have been limited by the lack of biologically plausible neuronal readout architectures; a limitation overcome by our model, which allows evaluation of specific components of the auditory pathway in a functional context.

Implications for Speech Processing

Psychoacoustic studies have indicated that the neural mechanism underlying the perceptual normalization of temporal speech cues is involuntary, i.e., it is cognitively impenetrable [16], controlled by physical rather than perceived speaking rate [17], confined to a temporally local context [2,18], not specific to speech sounds [52], and already operational in prearticulate infants [53]. The proposed conductance-based time-rescaling mechanism is consistent with these constraints. Moreover, our model posits a direct functional relation between high synaptic conductances and the time-warp robustness of human speech perception. This relation gives rise to a novel mechanistic hypothesis explaining the impaired capabilities of elderly listeners to process time-compressed speech [54,55]. We hypothesize that the down-regulation of inhibitory neurotransmitter systems in aging mammalian auditory pathways [56,57] limits the total synaptic conductance and therefore prevents the time-rescaling mechanism from generating short, effective time constants through synaptic shunting. Furthermore, our model implies that comprehension deficits in older adults should be linked specifically to the processing of phonetic segments that contain fast time-compressed temporal cues. Our hypothesis is consistent with two interrelated lines of evidence. First, comprehension difficulties of time-compressed speech in older adults are more likely a consequence of an age-related decline in central auditory processing than attributes of a general cognitive slowing [56,58]. Second, recent reports have indicated that recognition differences between young and elderly listeners originate mainly from the temporal compression of consonants, which often feature rapid spectral transitions, but not from steady-state segments [54,55,58] of speech. Finally, our hypothesis posits that speaking rate–induced shifts in perceptual category boundaries [2,16,17] should be age-dependent, i.e., their magnitude should decrease with increasing listener age. This prediction is straightforwardly testable within established psychoacoustic paradigms.

Connections to Other Models of Time-Warp–Invariant Processing

In a previous neuronal model of time-warp–invariant speech processing [5,6], sequences of acoustic events are converted into patterns of transiently matching firing rates in subsets of neurons within a population, which trigger synchronous firing in a downstream readout circuit. The identity of neurons whose firing rates converge to an identical value during an input pattern, and hence also the pattern of synchrony emerging in the readout layer, depends only on the relative timing of the events, not on the absolute duration of the auditory signal. However, for this model to recognize multiple input patterns, the convergence of firing rates is only approximate. Therefore, the resulting time-warp robustness is limited and, as in our model, dependent on the learning load. Testing this model on our synthetic classification task (cf. Figure 3) indicated a substantially smaller storage capacity than is realizable by the conductance-based tempotron (Text S1). An additional disadvantage of this approach is that it copes only with global (uniform) temporal warping. Invariant processing of dynamic time warp as is exhibited by natural speech (cf. Figure 1C and 1D) is more challenging since it requires robustness to local temporal distortions of a certain statistical character. Established algorithms that can cope with dynamically time-warped signals are typically based on minimizing the deviation between an observed signal and a stored reference template [59–61]. These algorithms are computationally expensive and lack biologically plausible neuronal implementations. By contrast, our conductance-based time-rescaling mechanism results naturally from the biophysical properties of input integration at the neuronal membrane and does not require dedicated computational resources. Importantly, our model does not rely on a comparison between the incoming signal and a stored reference template. Rather, after synaptic conductances have adjusted to the statistics of a given stimulus ensemble, the mechanism generalizes and automatically stabilizes neuronal voltage responses against dynamic time warp even when processing novel stimuli (cf. Figure 3C). The architecture of our neuronal model also fundamentally departs from the decades-old layout of Hidden Markov Model–based artificial speech-recognition systems, which employ probabilistic models of state sequences. These systems are hard to reconcile with the biological reality of neuronal system architecture, dynamics, and plasticity. The similarity in performance between our model and such state-of-the-art systems on a small vocabulary task highlights the powerful processing capabilities of spike-based neural representations and computation.

Generality of Mechanism

Although the present work focuses on the concrete and well-documented example of time-warp robustness in the context of neural speech processing, the proposed mechanism of automatic rescaling of integration time is general and applies also to other problems of neuronal temporal processing such as birdsong recognition [3], insect communication [9], and other ethologically important natural auditory signals. Moreover, robustness of neuronal processing to temporal distortions of spike patterns is not only important for the processing of stimulus time dependencies, but also in the context of spike-timing–based neuronal codes in which the precise temporal structure of spiking activity encodes

information about nontemporal physical stimulus dimensions [62]. Evidence for such temporal neural codes have been reported in the visual [63–65], auditory [66], and somatosensory [67], as well as the olfactory [68] pathways. As a result, we expect mechanisms of time-warp–invariant processing to also play a role in generating perceptual constancies along nontemporal stimulus dimensions such as contrast invariance in vision or concentration invariance in olfaction [4]. Finally, time warp has also been described in intrinsically generated brain signals. Specifically, the replay of hippocampal and cortical spiking activity at variable temporal warping [69,70] suggests that our model has applicability beyond sensory processing, possibly also encompassing memory storage and retrieval.

Materials and Methods

Conductance-Based Neuron Model

Numerical simulations of the conductance-based tempotron were based on exact integration [71] of the conductance-based voltage dynamics defined in Equation 1. With the membrane capacitance set to 1, the resting membrane time constant in this model is $\tau_m = 1/g_{leak}$. Implementing an integrate-and-fire neuron model, an output spike was elicited when $V(t)$ crossed the firing threshold V_{thr}. After a spike at t_{spike}, the voltage is smoothly reset to the resting value by shunting all synaptic inputs that arrive after t_{spike} (cf. [22]). We used $V_{thr} = 1$, $V_{rest} = 0$, and reversal potentials $V_{ex}^{rev} = 5$ and $V_{in}^{rev} = -1$ for excitatory and inhibitory conductances, respectively. Unless stated otherwise, the resting membrane time constant was set to $\tau_m = 100$ ms throughout our work [20]. For the synaptic time constant, we used $\tau_s = 1$ ms for the random latency task, which minimized the error of the current-based neuron, and to $\tau_s = 5$ ms in the speech-recognition tasks.

H-Current

To check the effect of nonsynaptic voltage-dependent conductances on the automatic time-rescaling mechanism, we implemented an h-like current I_h after [72] as a noninactivating current with HH-type dynamics of the form

$$I_h = g_h^{max} m\left(V - V_h^{rev}\right).$$

Here, g_h^{max} is the maximal h-conductance, with reversal potential $V_h^{rev} = -20$ mV and m is its voltage-dependent activation variable with kinetics

$$\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_h(V)}$$

where

$$\tau_h(V) = \frac{1}{\alpha(V) + \beta(V)}$$

and

$$m_\infty(V) = \frac{\alpha(V)}{\alpha(V) + \beta(V)}.$$

The voltage dependence of the rate constants α and β were described by the form

$$\alpha,\ \beta(V) = \frac{a_{\alpha,\beta} V + b_{\alpha,\beta}}{1 - \exp\left[\left(V + b_{\alpha,\beta}/a_{\alpha,\beta}\right)/k_{\alpha,\beta}\right]}$$

with parameters $a_\alpha = -39.015\ s^{-1}$, $b_\alpha = -259.925\ s^{-1}$, $k_\alpha = 1.77926$ and $a_\beta = 365.85\ s^{-1}$, $b_\beta = -2853.25\ s^{-1}$, $k_\beta = -1.28889$.

In Figure S1, we quantified the effect of the h-conductance on the fidelity of the time-rescaling mechanism by measuring the time-warp–induced distortions of neuronal voltage traces for different values of the maximal h-conductance g_h^{max}. Specifically, for a given value of g_h^{max} and a time warp β, we measure the voltage traces $V_{g_h^{max}}(t, 1)$ and $V_{g_h^{max}}(t, \beta)$ and their standard deviations across time σ_1 and σ_β, respectively. We define the time-warp distortion index $\Lambda\left(g_h^{max}, \beta\right)$ as the mean magnitude of the time-warp–induced voltage difference across time normalized by the mean standard deviation, $\bar{\sigma} = (\sigma_1 + \sigma_\beta)/2$,

$$\Lambda\left(g_h^{max}, \beta\right) = \frac{\left\langle \left| V_{g_h^{max}}(t, 1) - V_{g_h^{max}}(\beta t, \beta) \right| \right\rangle_t}{\bar{\sigma}}.$$

In Figure S1, values of $\Lambda\left(g_h^{max}, \beta\right)$ are normalized by $\Lambda(0,\beta)$. The voltage traces were generated by random latency patterns and uniform synaptic peak conductances as used in Figure 2. As increasing values of g_h^{max} depolarized the neuron's resting potential, excitatory and inhibitory synaptic conductances were balanced separately for each value of g_h^{max}.

Current-Based Neuron Model

In the current-based tempotron that was implemented as described in [22], each input spike evoked an exponentially decaying synaptic current that gave rise to a postsynaptic potential with a fixed temporal profile. In Figure 8C (upper row), voltage traces of a current-based analog of a conductance-based tempotron with learned synaptic conductances g_i^{max}, reversal potentials V_i^{rev}, and effective membrane integration time τ_{eff} (cf. Figure 8B) were computed by setting the synaptic efficacies ω_i of the current-based neuron to $\omega_i = g_i^{max} V_i^{rev}$ and its membrane time constant to $\tau_m = \tau_{eff}$. The resulting current-based voltage traces were scaled such that for each pair of models, the mean voltage maxima for unwarped stimuli ($\beta = 1$) were equal.

Tempotron Learning

Following [22], changes in the synaptic peak conductance g_i^{max} of the ith synapse after an error trial were given by the gradient of the postsynaptic potential, $\Delta g_i^{max} \propto -dV(t_{max})/dg_i^{max}$, at the time of its maximal value t_{max}. To compute the synaptic update for a given error trial, the exact solution of Equation 1 was differentiated with respect to g_i^{max} and evaluated at t_{max}, which was determined numerically for each error trial. Whenever a synaptic peak conductance attempted to cross to a negative value, its reversal potential was switched.

Voltage Correlation-Based Learning

A voltage correlation-based approximation of tempotron learning was implemented by extending the approach in [22] such that the change in the synaptic peak conductance g_i^{max} of the ith synapse due to a spike at time t_i was governed by the correlation $v_i = \int_{t_i}^{\infty} dt V(t) K_{learn}(t - t_i)$ of the postsynaptic potential $V(t)$ with a synaptic learning kernel $K_{learn}(t) = (\exp(-t/\tau_{learn}) - \exp(-t/\tau_s))/(\tau_{learn} - \tau_s)$. The two time constants of the

synaptic learning kernel were the synaptic time constant τ_s and the learning time constant $\tau_{\text{learn}} = 1/(g_{\text{leak}} + \overline{G}_{\text{syn}})$, which was determined by the time-averaged synaptic conductance $\overline{G}_{\text{syn}}$ of the current error trial and approximated the effective membrane time constant during that trial. The voltage maximum operation was approximated by thresholding v_i, yielding

$$\Delta g_i^{\max} \propto \begin{cases} \pm 1 & v_i > \kappa \\ 0 & v_i \leq \kappa \end{cases}$$

for changes of excitatory conductances on target and null patterns, respectively, and changes with the reversed polarity, ± 1, for inhibitory conductances. The plasticity induction threshold was set to $\kappa = 0.45$.

Learning Rate and Momentum Term

As in [22], we employed a momentum heuristic to accelerate learning in all learning rules. In this scheme, synaptic updates $[\Delta g_i^{\max}]_{\text{current}}$ consisted, not only of the correction $\lambda \Delta g_i^{\max}$, which was given by the learning rule and the learning rate λ, but also incorporated a fraction μ of the previous synaptic change $[\Delta g_i^{\max}]_{\text{previous}}$. Hence, $[\Delta g_i^{\max}]_{\text{current}} = \lambda \Delta g_i^{\max} + \mu [\Delta g^{\max}]_{\text{previous}}$. We used an adaptive learning rate that decreased from its initial value λ_{ini} as the number of learning cycles l grew, $\lambda = \lambda_{\text{ini}}/(1+10^{-4}(l-1))$. A learning cycle corresponded to one iteration through the batch of templates in the random latency task or the training set in the speech task.

Random latency task training. To ensure a fair comparison between the conductance-based and the current-based tempotrons (cf. Figure 3A), the learning rule parameters λ_{ini} and μ were optimized for each model. Specifically, for each value of $\beta_{\max}$, optimal values over a 2-dimensional grid were determined by the minimal error frequency achieved during runs over 10^5 cycles, with synaptic efficacies starting from Gaussian distributions with zero mean and standard deviations of 0.001. The optimization was performed over five realizations.

Global Time Warp

Global time warp was implemented by multiplying all firing times of a spike template by a constant scaling factor β. In Figure 3A, random global time warp between compression by $1/\beta_{\max}$ and dilation by $\beta_{\max}$ was generated by setting $\beta = \exp(q \ln(\beta_{\max}))$ with q drawn from a uniform distribution between -1 and 1 for each presentation.

Dynamic Time Warp

Dynamic time warp was implemented by scaling successive interspike intervals $t_j - t_{j-1}$ of a given template with a time-dependent warping factor $\bar{\beta}(t)$, such that warped spike times $t'_j = t'_{j-1} + \bar{\beta}(t_j)(t_j - t_{j-1})$ with $t'_1 \equiv t_1$ and $\bar{\beta}(t) = \exp(\tilde{q}(t) \ln(\beta_{\max}))$. The time-dependent factor $\tilde{q}(t) = \text{erfc}(\xi(t)) - 1$ resulted from an equilibrated Ornstein-Uhlenbeck process $\xi(t)$ with a relaxation time of $\tau = 200$ ms that was rescaled by the complementary error function erfc to transform the normal distribution of $\xi(t)$ into a uniform distribution over $[-1\ 1]$ at each t.

Global Inhibition Model

To ensure that the symmetry of excitation and inhibition in our model architecture was not crucial for the time-warp–invariant processing of spike patterns, we implemented a control architecture in which all afferents were confined to be excitatory, except one additional inhibitory synapse, which mimicked the effect of a global inhibitory network. In the time-warped random latency task, spike patterns were fed into the excitatory population as before; however, in addition, the inhibitory synapse received a copy of each incoming spike. All synaptic peak conductances were plastic and controlled by the conductance-based tempotron rule. In this model, synaptic sign changes were prohibited.

Gaussian Spike Time Jitter

Spike time jitter [22] was implemented by adding independent Gaussian noise with zero mean and a standard deviation of 5 ms to each spike of a template before each presentation.

Acoustic Front-End

Sound signals were normalized to unit peak amplitude and converted into spectrograms over $N_{\text{FTT}} = 129$ linearly spaced frequencies $f_j = f_{\min} + j(f_{\max} + f_{\min})/(N_{\text{FTT}}+1)$ $(j = 1 \ldots N_{\text{FTT}})$ between $f_{\min} = 130$ Hz and $f_{\max} = 5{,}400$ Hz by a sliding fast Fourier transform with a window size of 256 samples and a temporal step size of 1 ms. The resulting spectrograms were filtered into $N_f = 32$ logarithmically spaced Mel frequency channels by overlapping triangular frequency kernels. Specifically, N_f+2 linearly spaced frequencies given by $h_j = h_{\min} + j(h_{\max} - h_{\min})/(N_f+1)$ with $j = 0 \ldots N_f+1$ and $h_{\max,\min} = 2{,}595 \log(1 + f_{\max,\min}/700)$ were transformed to a Mel frequency scale $f_j^{\text{Mel}} = 700(\exp(h_j/2595) - 1)$ between $f_{\min}$ and $f_{\max}$. Based on these, signals in N_f channels resulted from triangular frequency filters over intervals $\left[f_{j-1}^{\text{Mel}}, f_{j+1}^{\text{Mel}}\right]$ with center peaks at f_j^{Mel} $(j = 1 \ldots N_f)$. After normalization of the resulting Mel spectrogram S^{Mel} to unit peak amplitude, the logarithm was taken through $\log(S^{\text{Mel}} = \varepsilon) - \log(\varepsilon)$ with $\varepsilon = 10^{-5}$ and the signal in each frequency channel smoothed in time by a Gaussian kernel with a time constant of 10 ms. Spikes were generated by thresholding of the resulting signals by a total of 31 onset and offset threshold-crossing detector units. Whereas each onset afferent emitted a spike whenever the signal crossed its threshold in the upward direction, offset afferents fired when the signal dropped below the threshold from above. For each frequency channel and each utterance, threshold levels for onset and offset afferents were set relative to the maximum signal over time to $\vartheta_1 = 0.01$ and $\vartheta_j = j/15$ $(j = 1 \ldots 15)$. For $\vartheta_{15} = 1$, onset and offset afferents were reduced to a single afferent whose spikes encoded the time of the maximum signal for a given frequency channel.

Speech Databases

We used the digit subset of the TI46 Word speech database [24]. This clear speech dataset comprises 26 isolated utterances of each English digit from zero to nine spoken by 16 adult speakers (eight male and eight female). The data is partitioned into a fixed training set, comprising 10 utterances per digit and speaker, and a fixed test set holding the remaining 16 utterances per digit and speaker. We also tested our neuronal word-recognition model on the adult speaker, isolated-digit subset of the TIDIGITS database [27]. This subset comprises two utterances per digit and speaker, i.e., a total of 20 utterances from 225 adult speakers (111 male and 114 female), that are dialectically balanced across 21 dialectical regions (tiling the continental United States). Because the TI46 database only provides utterances of the word "zero" for the digit 0, we excluded the utterances of "oh" from our TIDIGITS sample.

Digit Classification

Based on the spiking activity of all binary digit detector neurons, a full digit classifier was implemented by ranking the digit

detectors according to their individual task performances. As a result, a given stimulus was classified as the target digit of the most reliable of all responding digit detector neurons. If all neurons remained silent, a stimulus was classified as the target digit of the least reliable neuron.

Spike-Triggered Target Features

To preserve the timing relations between the learned spectro-temporal features and the target words, we refrained from correcting the spike-triggered stimuli for stimulus autocorrelations [73].

Speech Task Training

Test errors in the speech tasks were substantially reduced by training with a Gaussian spike jitter with a standard deviation of σ added to the input spikes as well a symmetric threshold margin v that required the maximum postsynaptic voltage on target stimuli to exceed $V_{\text{thr}}+v$ and to remain below $V_{\text{thr}}-v$ during null stimuli. Values of λ_{ini}, μ, σ, and v were optimized on a 4-dimensional grid. Because for each grid point, only short runs over maximally 200 cycles were performed, we also varied the mean values of initial Gaussian distributions of the excitatory and inhibitory synaptic peak conductances, keeping their standard deviations fixed at 0.001. The reported performances are based on the solutions that had the smallest errors fractions over the test set. If not unique, we selected the solution with the highest robustness to time warp (cf. Figure 8B). Note that this naive optimization of the training parameters did not maintain a separate holdout test set for cross-validation and might therefore overestimate the true generalization performance. We adopted this optimization scheme from [25,26] to ensure comparability of the resulting performance measures.

Comparison to the HTK

HTK generalization performance was tested with the HTK package version 3.4.1 [74] with front-end and HMM model parameters following [26]. Specifically, speech data from the TI46 and TIDIGITS databases were converted to 13 Mel-cepstral coefficients (including the 0*th* order coefficient) along with their first and second derivatives at a frame rate of 5 ms. Mel-coefficients were computed over 30 channels in 25-ms windows with zero mean normalization enabled (TARGETKIND = MFCC_D_A_Z_0). In addition, the following parameters were set: USEHAMMING = T; PREEMPCOEF = 0.97; and CEPLIFTER = 22. Ten HMM models, one for each digit plus one HMM model for silence, were used. Each model consisted of five states (including the the two terminal states) with eight Gaussian mixtures per state and left-to-right (no skip) transition topology.

Supporting Information

Figure S1 Effect of h-conductance on time rescaling. Time-warp distortion index computed for random latency patterns (see Materials and Methods) versus the maximal h-conductance for different values of the mean synaptic conductance $\overline{G}_{\text{syn}}/g_{\text{leak}}$: 7.2 (triangles), 10.8 (squares), and 14.4 (circles). Curves were averaged over 2,000 spike-pattern realizations.
Found at: doi:10.1371/journal.pbio.1000141.s001 (0.70 MB TIF)

Table S1 Generalization from TI46 to TIDIGITS. For each dialect group, the table lists the percentages of speakers for which our model committed a given number of word-recognition errors.
Found at: doi:10.1371/journal.pbio.1000141.s002 (0.01 MB PDF)

Text S1 Comparison to the Hopfield-Brody model of time-warp–invariant neuronal processing.
Found at: doi:10.1371/journal.pbio.1000141.s003 (0.03 MB PDF)

Acknowledgments

We thank C. Brody, D. Buonomano, D. Hansel, M. Kilgard, M. London, M. Merzenich, J. Miller, I. Nelken, A. Roth, S. Shamma, and M. Tsodyks for discussions and comments.

Author Contributions

The author(s) have made the following declarations about their contributions: Conceived and designed the experiments: RG HS. Performed the experiments: RG HS. Analyzed the data: RG HS. Contributed reagents/materials/analysis tools: RG HS. Wrote the paper: RG HS.

References

1. Sakoe H, Chiba S (1978) Dynamic programming algorithm optimization for spoken word recognition. IEEE Trans Acoust Speech Signal Process 26: 43–49.
2. Miller JL (1981) Effects of speaking rate on segmental distinctions. In: Eimas PD, Miller JL, eds. Perspectives on the study of speech. Hilsdale (New Jersey): Lawrence Erlbaum Associates. pp 39–74.
3. Anderson S, Dave A, Margoliash D (1996) Template-based automatic recognition of birdsong syllables from continuous recordings. J Acoust Soc Am 100: 1209–1219.
4. Hopfield JJ (1996) Transforming neural computations and representing time. Proc Natl Acad Sci U S A 93: 15440–15444.
5. Hopfield JJ, Brody CD (2000) What is a moment? "Cortical" sensory integration over a brief interval. Proc Natl Acad Sci U S A 97: 13919–131924.
6. Hopfield JJ, Brody CD (2001) What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. Proc Natl Acad Sci U S A 98: 1282–1287.
7. Herz AVM (2005) How is time represented in the brain? In: van Hemmen J, Sejnowski T, eds. 23 problems in systems neuroscience. Oxford (United Kingdom): Oxford University Press. pp 266–283.
8. Brown J, Miller P (2007) Automatic classification of killer whale vocalizations using dynamic time warping. J Acoust Soc Am 122: 1201–1207.
9. Gollisch T (2008) Time-warp invariant pattern detection with bursting neurons. New J Phys 10: 015012.
10. Shannon R, Zeng F, Kamath V, Wygonski J, Ekelid M (1995) Speech recognition with primarily temporal cues. Science 270: 303–304.
11. Merzenich M, Jenkins W, Johnston P, Schreiner C, Miller S, et al. (1996) Temporal processing deficits of language-learning impaired children ameliorated by training. Science 271: 77–81.
12. Phillips D, Farmer M (1990) Acquired word deafness, and the temporal grain of sound representation in the primary auditory cortex. Behav Brain Res 40: 85–94.
13. Fitch RH, Miller S, Tallal P (1997) Neurobiology of speech perception. Annu Rev Neurosci 20: 331–351.
14. Miller JL, Grosjean F, Lomanto C (1984) Articulation rate and its variability in spontaneous speech: a reanalysis and some implications. Phonetica 41: 215–225.
15. Miller JL, Grosjean F, Lomanto C (1986) Speaking rate and segments: a look at the relation between speech production and speech perception for voicing contrast. Phonetica 43: 106–115.
16. Miller JL, Green K, Schermer TM (1984) A distinction between the effects of sentential speaking rate and semantic congruity on word identification. Percept Psychophys 36: 329–337.
17. Miller JL, Aibel IL, Green K (1984) On the nature of rate-dependent processing during phonetic perception. Percept Psychophys 35: 5–15.
18. Newman R, Sawusch J (1996) Perceptual normalization for speaking rate: effects of temporal distance. Percept Psychophys 58: 540–560.
19. Bernander O, Douglas R, Martin K, Koch C (1991) Synaptic background activity influences spatiotemporal integration in single pyramidal cells. Proc Natl Acad Sci U S A 88: 11569–11573.
20. Koch C, Rapp M, Segev I (1996) A brief history of time (constants). Cereb Cortex 6: 93–101.
21. Häusser M, Clark BA (1997) Tonic synaptic inhibition modulates neuronal output pattern and spatiotemporal synaptic integration. Neuron 19: 665–678.
22. Gütig R, Sompolinsky H (2006) The tempotron: a neuron that learns spike timing-based decisions. Nat Neurosci 9: 420–428.

23. Hopfield JJ (2004) Encoding for computation: recognizing brief dynamical patterns by exploiting effects of weak rhythms on action-potential timing. Proc Natl Acad Sci U S A 101: 6255–6260.
24. Liberman M, Amsler R, Church K, Fox E, Hafner C, et al. (1993) TI 46-Word. Philadelphia (Pennsylvania): Linguistic Data Consortium.
25. Walker W, Lamere P, Kwok P, Raj B, Singh R, et al. (2004) Sphinx-4: a flexible open source framework for speech recognition. Technical Report SMLI TR-2004-139. Menlo Park (California): Sun Microsystems Laboratories. pp 1–15.
26. Deshmukh O, Espy-Wilson C, Juneja A (2002) Acoustic-phonetic speech parameters for speaker-independent speech recognition. In: Proceedings of IEEE ICASSP 2002; 13–17 May 2002. Orlando, Florida, United States. pp 593–596.
27. Leonard R, Doddington G (1993) TIDIGITS. Philadelphia (Pennsylvania): Linguistic Data Consortium.
28. Destexhe A, Rudolph M, Paré D (2003) The high-conductance state of neocortical neurons in vivo. Nat Rev Neurosci 4: 739–751.
29. Zhang L, Tan A, Schreiner C, Merzenich M (2003) Topography and synaptic shaping of direction selectivity in primary auditory cortex. Nature 424: 201–205.
30. Wehr M, Zador A (2003) Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. Nature 426: 442–446.
31. Borg-Graham L, Monier C, Frégnac Y (1998) Visual input evokes transient and strong shunting inhibition in visual cortical neurons. Nature 393: 369–373.
32. Hirsch JA, Alonso JM, Reid R, Martinez L (1998) Synaptic integration in striate cortical simple cells. J Neurosci 18: 9517–9528.
33. Shu Y, Hasenstaub A, McCormick DA (2003) Turning on and off recurrent balanced cortical activity. Nature 423: 288–293.
34. Haider B, Duque A, Hasenstaub AR, McCormick DA (2006) Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. J Neurosci 26: 4535–4545.
35. Waters J, Helmchen F (2006) Background synaptic activity is sparse in neocortex. J Neurosci 26: 8267–8277.
36. Major G, Larkman A, Jonas P, Sakmann B, Jack J (1994) Detailed passive cable models of whole-cell recorded ca3 pyramidal neurons in rat hippocampal slices. J Neurosci 14: 4613–4638.
37. Roth A, Häusser M (2001) Compartmental models of rat cerebellar purkinje cells based on simultaneous somatic and dendritic patch-clamp recordings. J Physiol 535: 445–572.
38. Sarid L, Bruno R, Sakmann B, Segev I, Feldmeyer D (2007) Modeling a layer 4-to-layer 2/3 module of a single column in rat neocortex: interweaving in vitro and in vivo experimental observations. Proc Natl Acad Sci U S A 104: 16353–16358.
39. Oswald A, Reyes A (2008) Maturation of intrinsic and synaptic properties of layer 2/3 pyramidal neurons in mouse auditory cortex. J Neurophysiol 99: 2998–3008.
40. Froemke R, Merzenich M, Schreiner C (2007) A synaptic memory trace for cortical receptive field plasticity. Nature 450: 425–429.
41. Froemke R, Dan Y (2002) Spike-timing-dependent synaptic modification induced by natural spike trains. Nature 416: 433–438.
42. Wang HX, Gerkin RC, Nauen DW, Bi GQ (2005) Coactivation and timing-dependent integration of synaptic potentiation and depression. Nat Neurosci 8: 187–193.
43. Froemke R, Tsay I, Raad M, Long J, Dan Y (2006) Contribution of individual spikes in burstinduced long-term synaptic modification. J Neurophysiol 95: 1620–1629.
44. Wittenberg G, Wang S (2006) Malleability of spike-timing-dependent plasticity at the ca3-ca1 synapse. J Neurosci 26: 6610–6617.
45. Zatorre R, Belin P (2001) Spectral and temporal processing in human auditory cortex. Cereb Cortex 11: 946–953.
46. Boemio A, Fromm S, Braun A, Poeppel D (2005) Hierarchical and asymmetric temporal sensitivity in human auditory cortices. Nat Neurosci 8: 389–395.
47. Abrams D, Nicol T, Zecker S, Kraus N (2008) Right-hemisphere auditory cortex is dominant for coding syllable patterns in speech. J Neurosci 28: 3958–3965.
48. Hutsler J, Galuske R (2003) Hemispheric asymmetries in cerebral cortical networks. Trends Neurosci 26: 429–435.
49. Shtyrov Y, Kujala T, Ahveninen J, Tervaniemi M, Alku P, et al. (1998) Background acoustic noise and the hemispheric lateralization of speech processing in the human brain: magnetic mismatch negativity study. Neurosci Lett 251: 141–144.
50. Abrams DA, Nicol T, Zecker SG, Kraus N (2006) Auditory brainstem timing predicts cerebral asymmetry for speech. J Neurosci 26: 11131–11137.
51. Casseday JH, Fremouw T, Covey E (2002) The inferior colliculus: a hub for the central auditory system. In: Oertel D, Fay R, Popper A, eds. Integrative functions in the mammalian auditory pathway. New York (New York): Springer. pp 238–318.
52. Jusczyk P, Pisoni D, Reed M, Fernald A, Myers M (1983) Infants' discrimination of the duration of a rapid spectrum change in nonspeech signals. Science 222: 175–177.
53. Eimas PD, Miller JL (1980) Contextual effects in infant speech perception. Science 209: 1140–1141.
54. Gordon-Salant S, Fitzgibbons P (2001) Sources of age-related recognition difficulty for timecompressed speech. J Speech Lang Hear Res 44: 709–719.
55. Gordon-Salant S, Fitzgibbons P, Friedman S (2007) Recognition of time-compressed and natural speech with selective temporal enhancements by young and elderly listeners. J Speech Lang Hear Res 50: 1181–1193.
56. Caspary D, Schatteman T, Hughes L (2005) Age-related changes in the inhibitory response properties of dorsal cochlear nucleus output neurons: role of inhibitory inputs. J Neurosci 25: 10952–10959.
57. Caspary DM, Ling L, Turner JG, Hughes LF (2008) Inhibitory neurotransmission, plasticity and aging in the mammalian central auditory system. J Exp Biol 211(Pt 11): 1781–1791.
58. Schneider BA, Daneman M, Murphy DR (2005) Speech comprehension difficulties in older adults: cognitive slowing or age-related changes in hearing? Psychol Aging 20: 261–271.
59. Itakura F (1975) Minimum prediction residual principle applied to speech recognition. IEEE Trans Acoust Speech Signal Proc ASSP-23: 67–72.
60. Myers C, Rabiner L, Rosenberg A (1980) Performance tradeoffs in dynamic time warping algorithms for isolated word recognition. IEEE Acoust Speech Signal Process ASSP-28: 623–635.
61. Kavaler RA, Lowy M, Murveit H, Brodersen RW (1987) A dynamic-time-warp integrated circuit for a 1000-word speech recognition system. IEEE J Solid-State Circuits 22: 3–14.
62. Mauk M, Buonomano D (2004) The neural basis of temporal processing. Annu Rev Neurosci 27: 307–340.
63. Meister M, Lagnado L, Baylor DA (1995) Concerted signaling by retinal ganglion cells. Science 270: 1207–1210.
64. Neuenschwander S, Singer W (1996) Long-range synchronization of oscillatory light responses in the cat retina and lateral geniculate nucleus. Nature 379: 728–732.
65. Gollisch T, Meister M (2008) Rapid neural coding in the retina with relative spike latencies. Science 319: 1108–1111.
66. deCharms RC, Merzenich MM (1996) Primary cortical representation of sounds by the coordination of action-potential timing. Nature 381: 610–613.
67. Johansson RS, Birznieks I (2004) First spikes in ensembles of human tactile afferents code complex spatial fingertip events. Nat Neurosci 7: 170–177.
68. Wehr M, Laurent G (1996) Odour encoding by temporal sequences of firing in oscillating neural assemblies. Nature 384: 162–166.
69. Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. Neuron 29: 145–156.
70. Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat Neurosci 10: 100–107.
71. Brette R (2006) Exact simulation of integrate-and-fire models with synaptic conductances. Neural Computat 18: 2004–2027.
72. Dickson CT, Magistretti J, Shalinsky MH, Fransén E, Hasselmo ME, et al. (2000) Properties and role of Ih in the pacing of subthreshold oscillations in entorhinal cortex layer II neurons. J Neurophysiol 83: 2562–2579.
73. Klein DJ, Depireux DA, Simon JZ, Shamma SA (2000) Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. J Comput Neurosci 9: 85–111.
74. Woodland P (2009) Htk3. Available: http://htk.eng.cam.ac.uk. Accessed 1 April 2009.

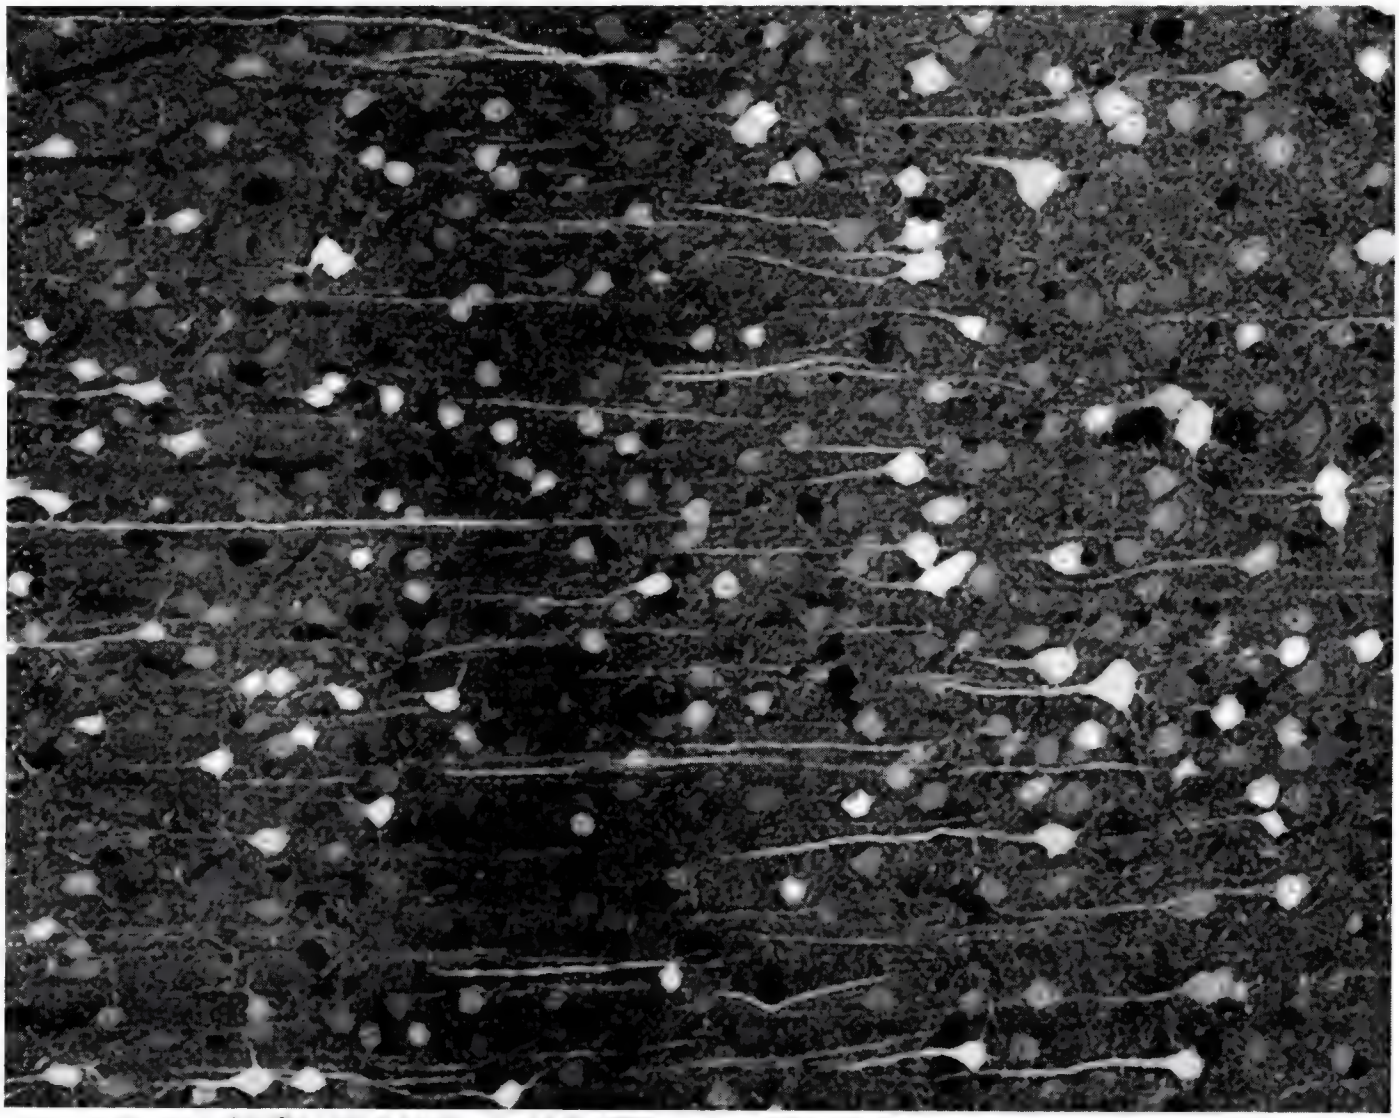

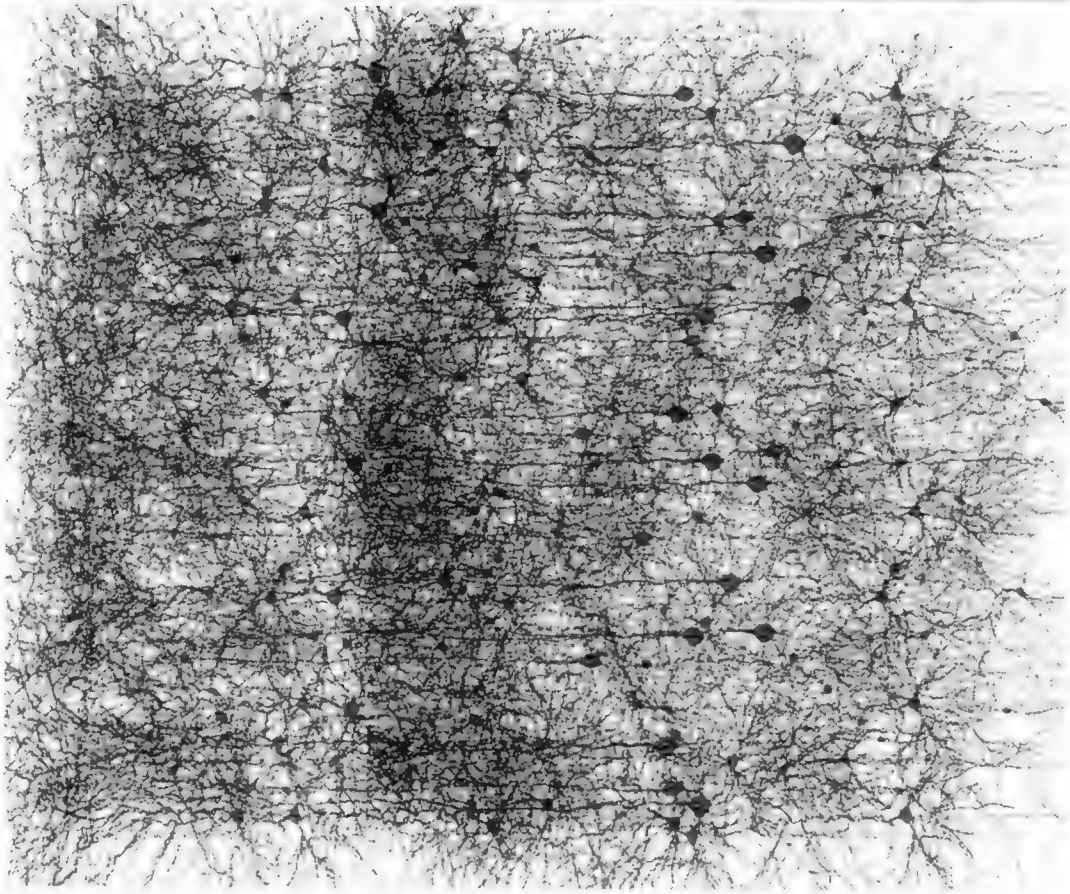

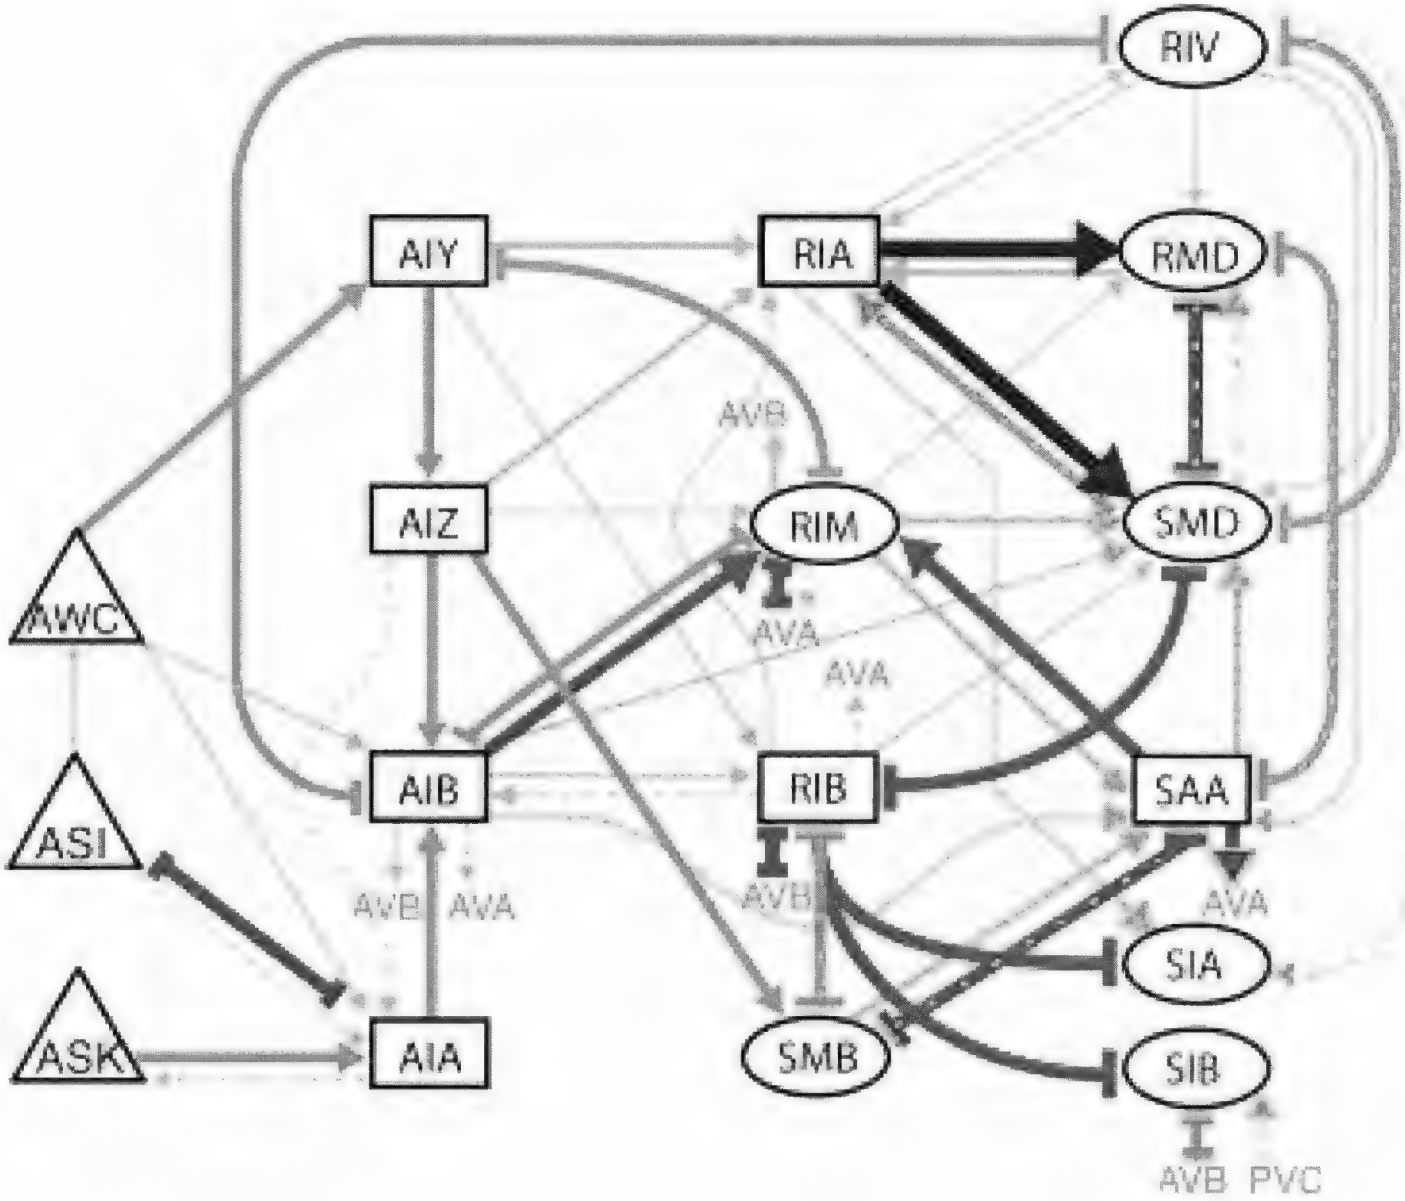

Gray, Hill, Bargmann 2005

Intelligent Design

Intelligent Design

simplify model

Intelligent Design

simplify model

match some anatomical and physiological data

Intelligent Design

simplify model

match some anatomical and physiological data

account for some behavior

Intelligent Design

simplify model

match some anatomical and physiological data

account for some behavior
(hopefully with efficient and "optimal" mechanism)

Intelligent Design

simplify model

match some anatomical and physiological data

account for some behavior
(hopefully with efficient and "optimal" mechanism)

complicate the model to match more data
and account for more behaviors

Intelligent Design

simplify model

match some anatomical and physiological data

account for some behavior
(hopefully with efficient and "optimal" mechanism)

complicate the model to match more data ← why?
and account for more behaviors

Intelligent Design

simplify model

match some anatomical and physiological data

account for some behavior
(hopefully with efficient and "optimal" mechanism)

complicate the model to match more data ← why?
and account for more behaviors ← ?!!?

Evolution

Evolution

start with a circuit that does something

Evolution

start with a circuit that does something

fiddle with it but do not destroy it

Evolution

start with a circuit that does something

fiddle with it but do not destroy it

hijack it to a new purpose

Evolution

start with a circuit that does something

fiddle with it but do not destroy it

hijack it to a new purpose

exploit any fortuitous interactions that arise

"Neural" Networks

Maass, Jaeger

"Neural" Networks

start with a circuit that is as complex as possible
and that you admittedly don't understand

Maass, Jaeger

"Neural" Networks

start with a circuit that is as complex as possible
and that you admittedly don't understand

fiddle with it but do not intrinsically change it

Maass, Jaeger

"Neural" Networks

start with a circuit that is as complex as possible
and that you admittedly don't understand

fiddle with it but do not intrinsically change it

hijack it to a useful purpose

Maass, Jaeger

$$\frac{dx_i}{dt} = -x_i + \sum_{j=1}^{N} J_{ij} r_j - b_i$$

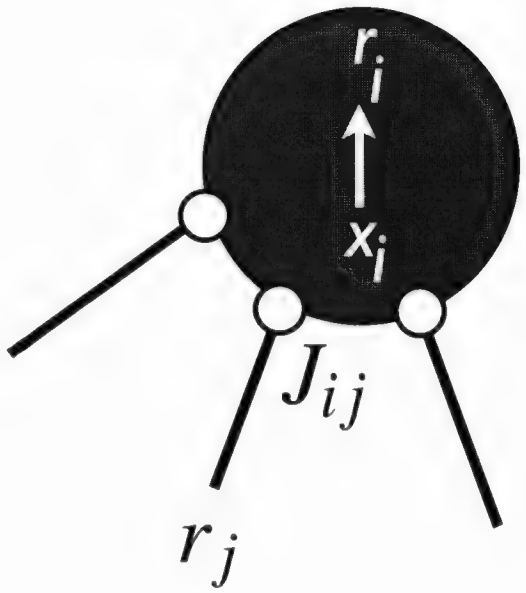

$$\frac{dx_i}{dt} = -x_i + \sum_{j=1}^{N} J_{ij} r_j - b_i$$

firing rate

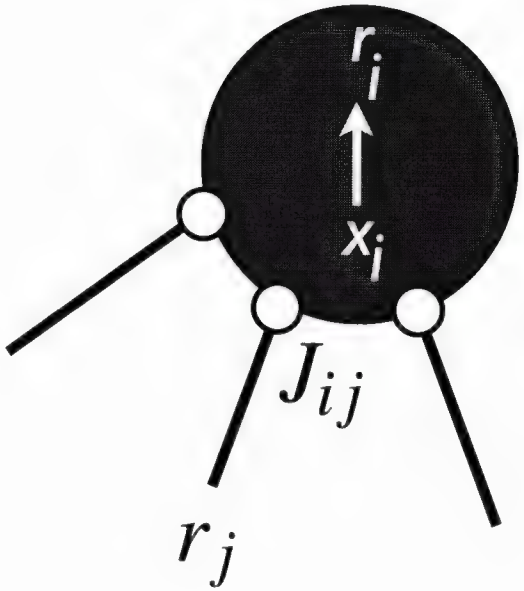

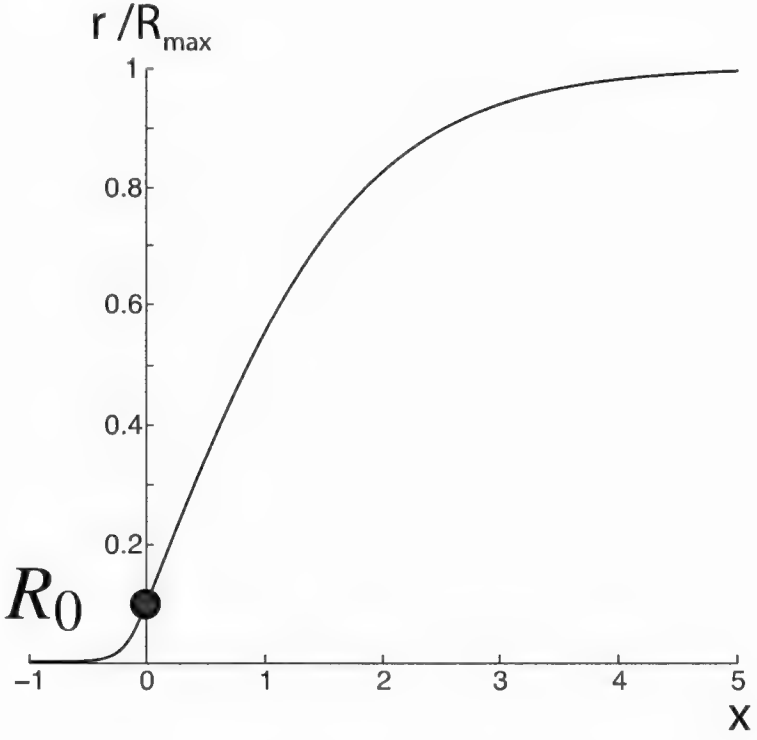

$$\frac{dx_i}{dt} = -x_i + \sum_{j=1}^{N} J_{ij} r_j - b_i$$

synaptic weights

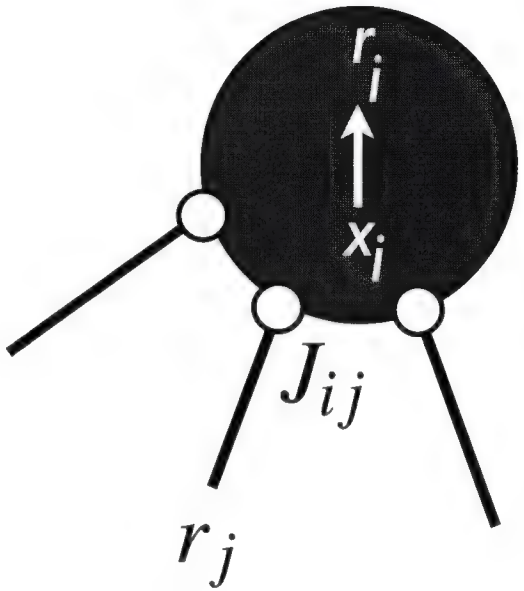

$$\langle J_{ij} \rangle_J = 0 \quad \text{and} \quad \langle J_{ij}^2 \rangle_J = \frac{g^2}{N}$$

$$\frac{dx_i}{dt} = -x_i + \sum_{j=1}^{N} J_{ij} r_j - b_i$$

bias $b_i = \sum J_{ij} R_0$

synaptic weights

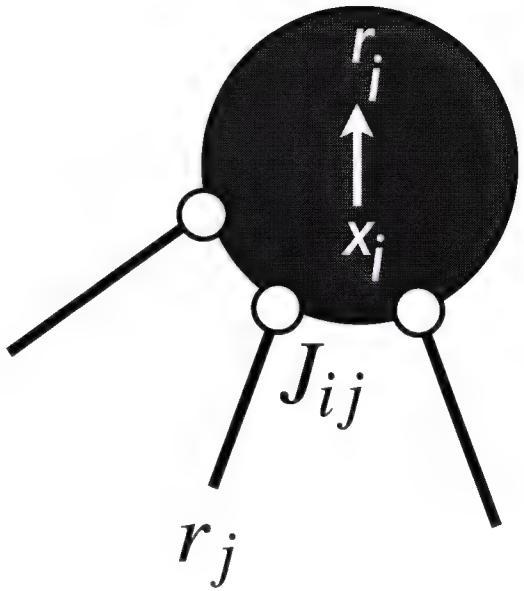

$$\langle J_{ij} \rangle_J = 0 \quad \text{and} \quad \langle J_{ij}^2 \rangle_J = \frac{g^2}{N}$$

$$\frac{dx_i}{dt} = -x_i + \sum_{j=1}^{N} J_{ij} r_j - b_i$$

bias $b_i = \sum J_{ij} R_0$

synaptic weights

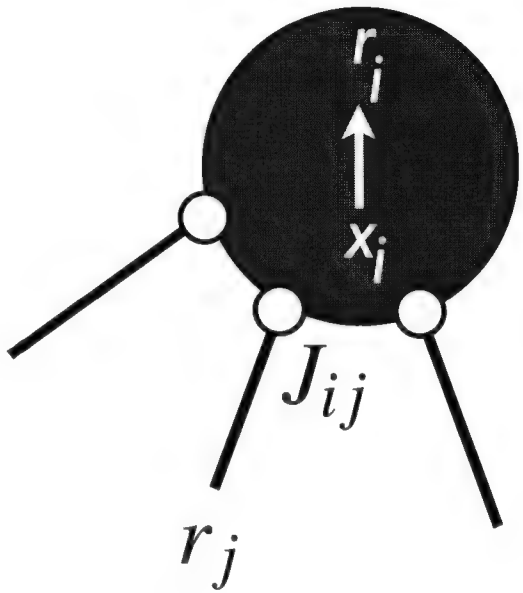

$$\langle J_{ij} \rangle_J = 0 \quad \text{and} \quad \langle J_{ij}^2 \rangle_J = \frac{g^2}{N}$$

When spontaneously active, these networks are chaotic

Sompolinsky, Crisanti & Sommers, 1988

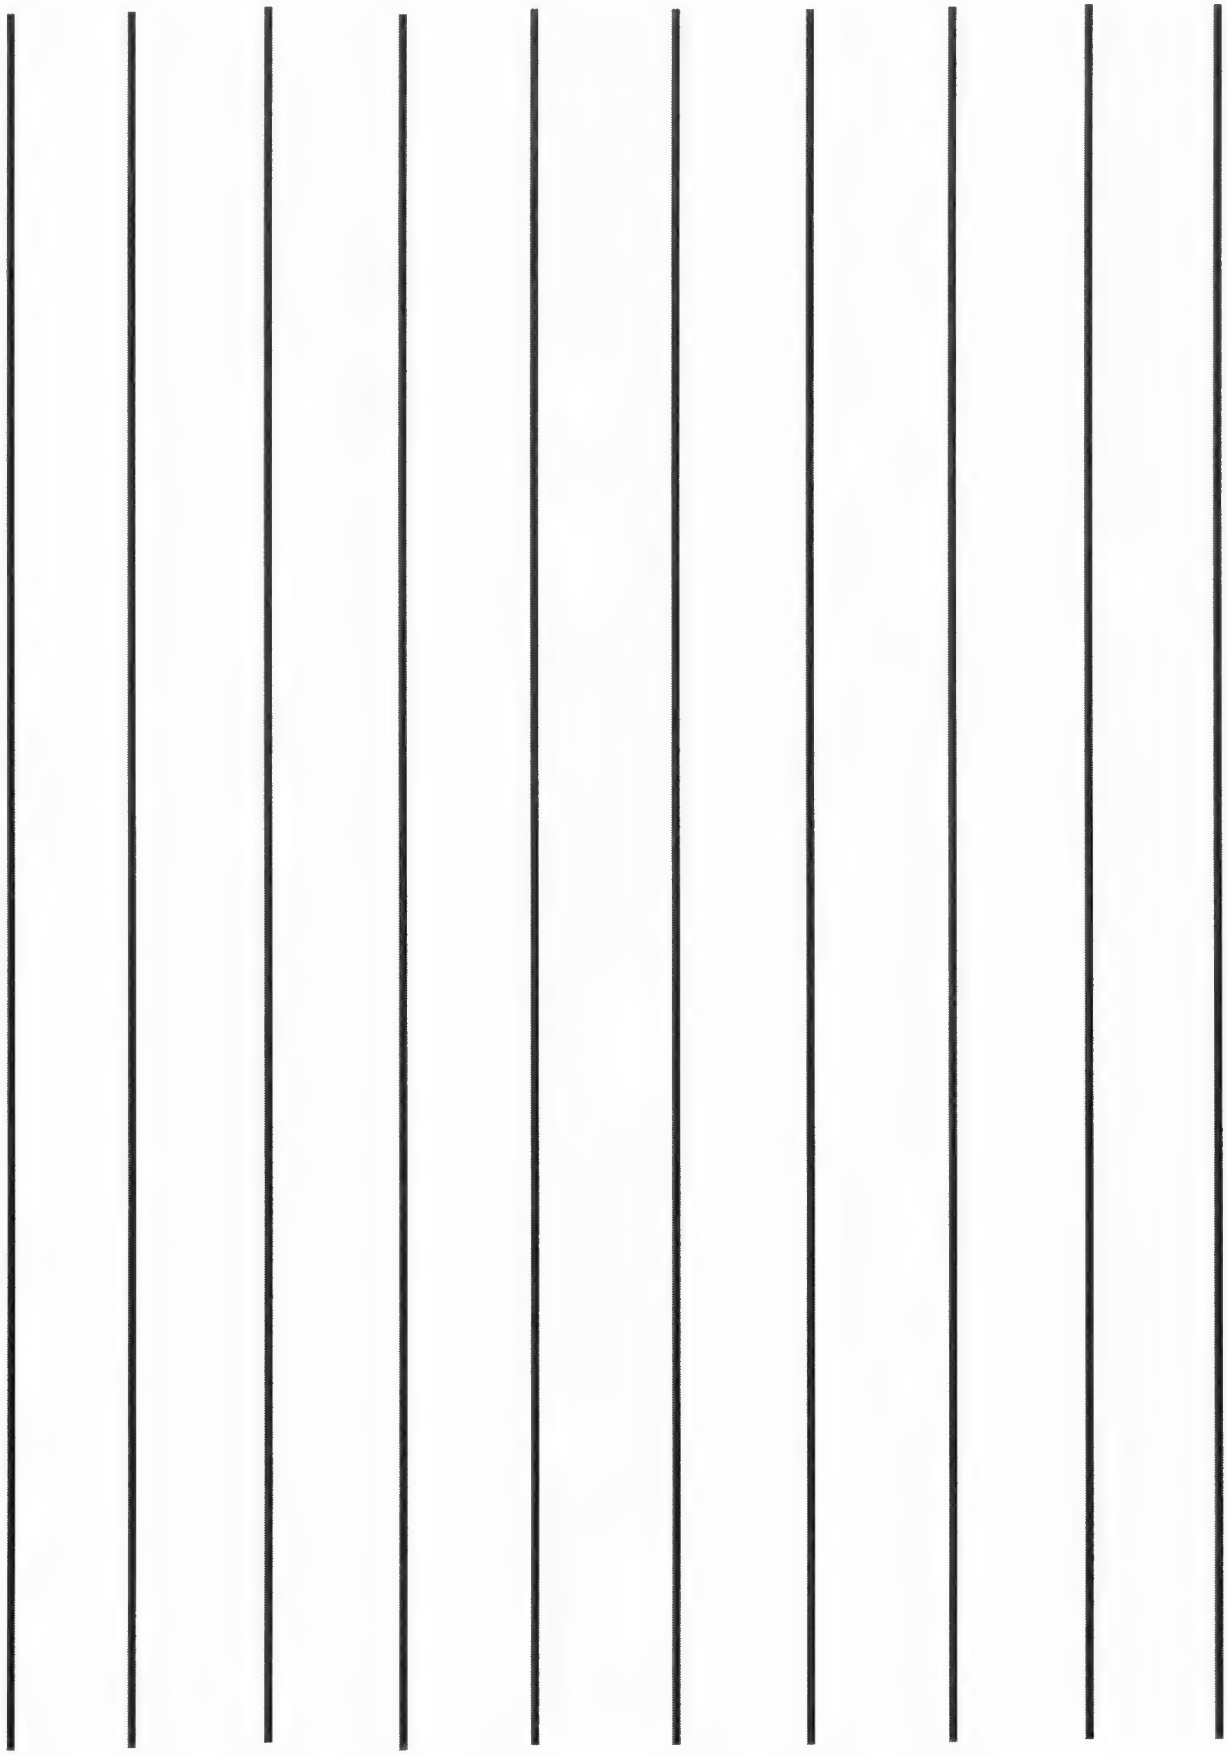

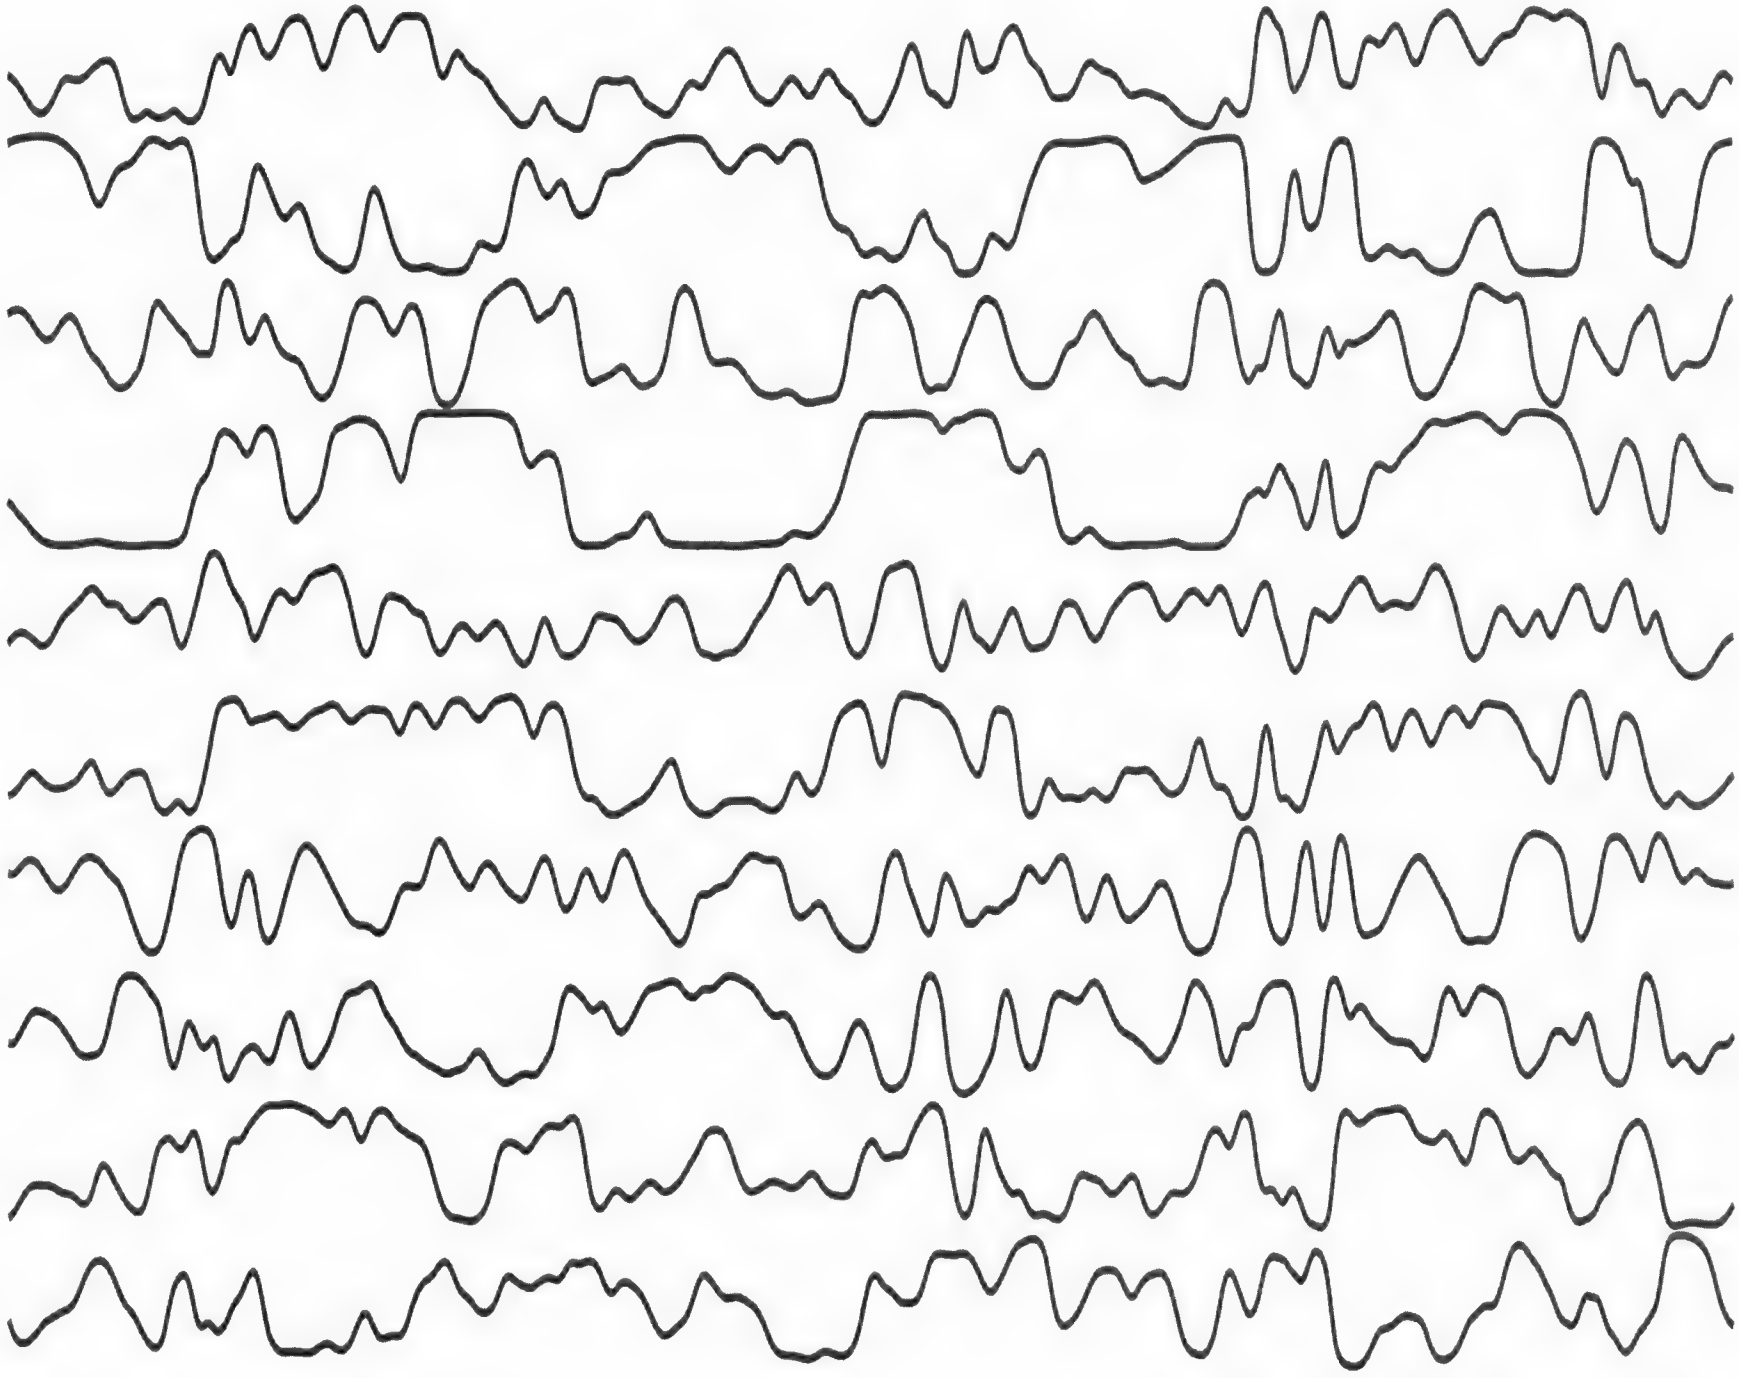

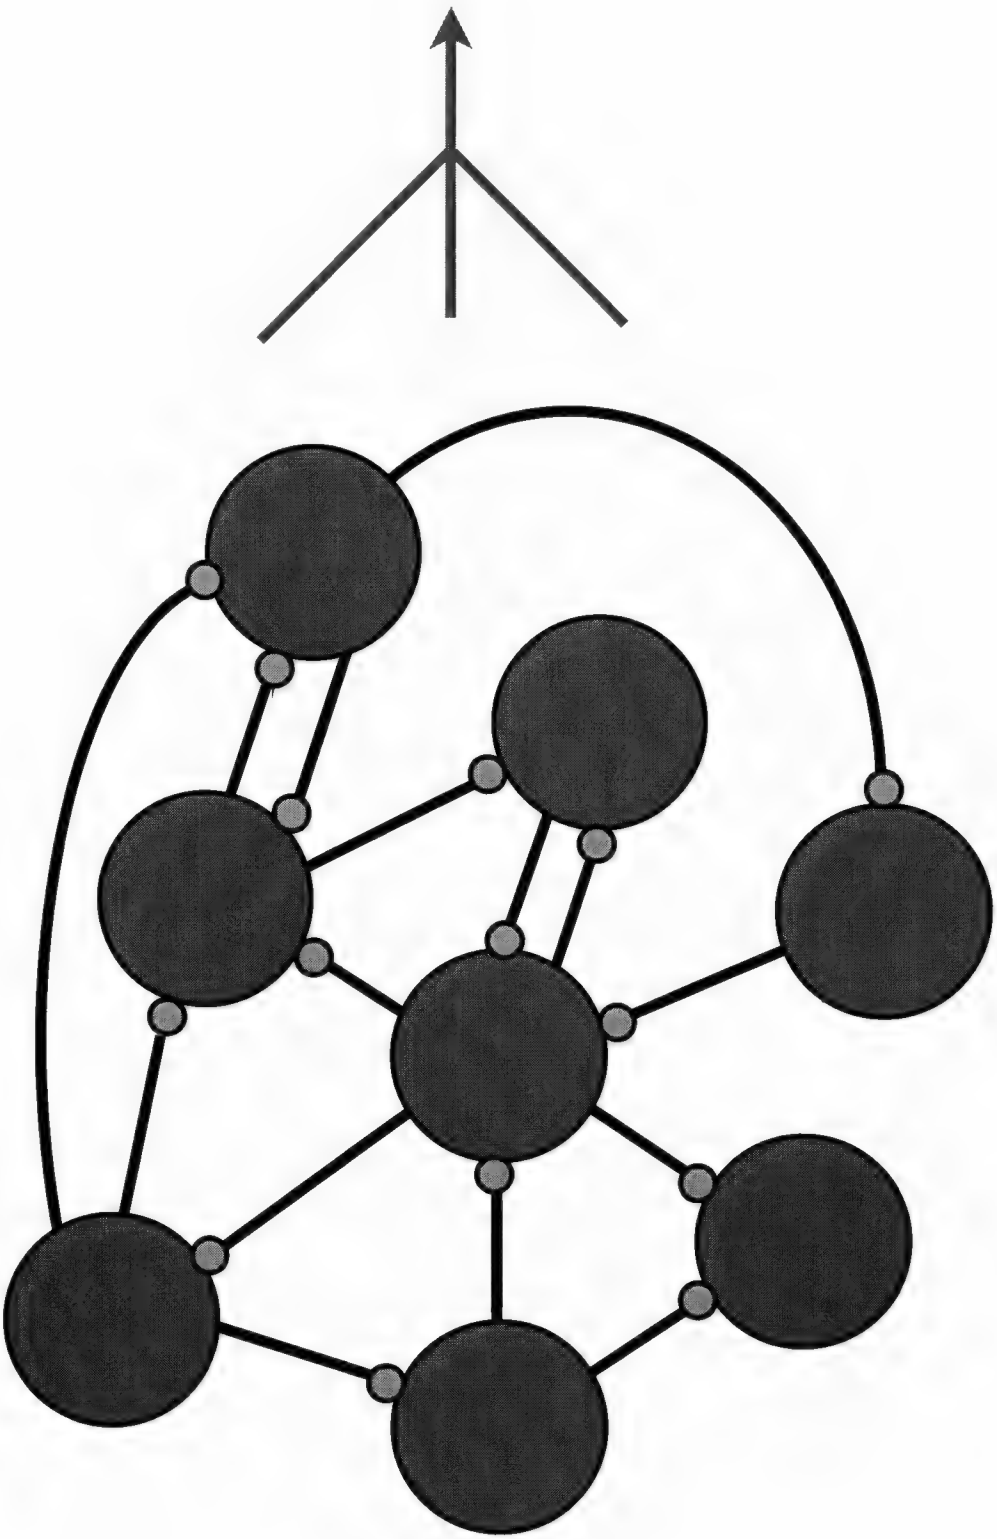

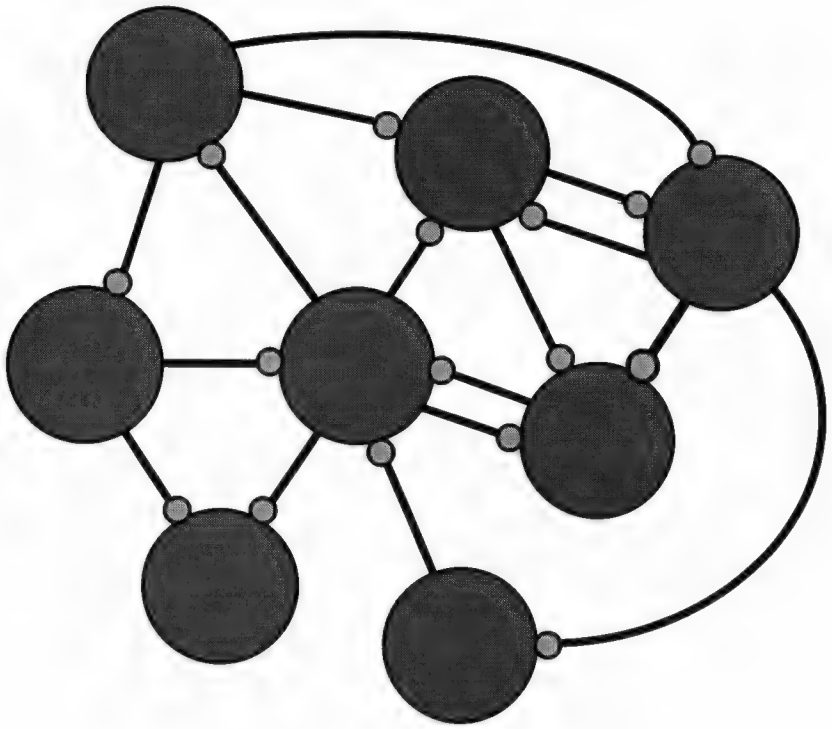

Maass, Jaeger

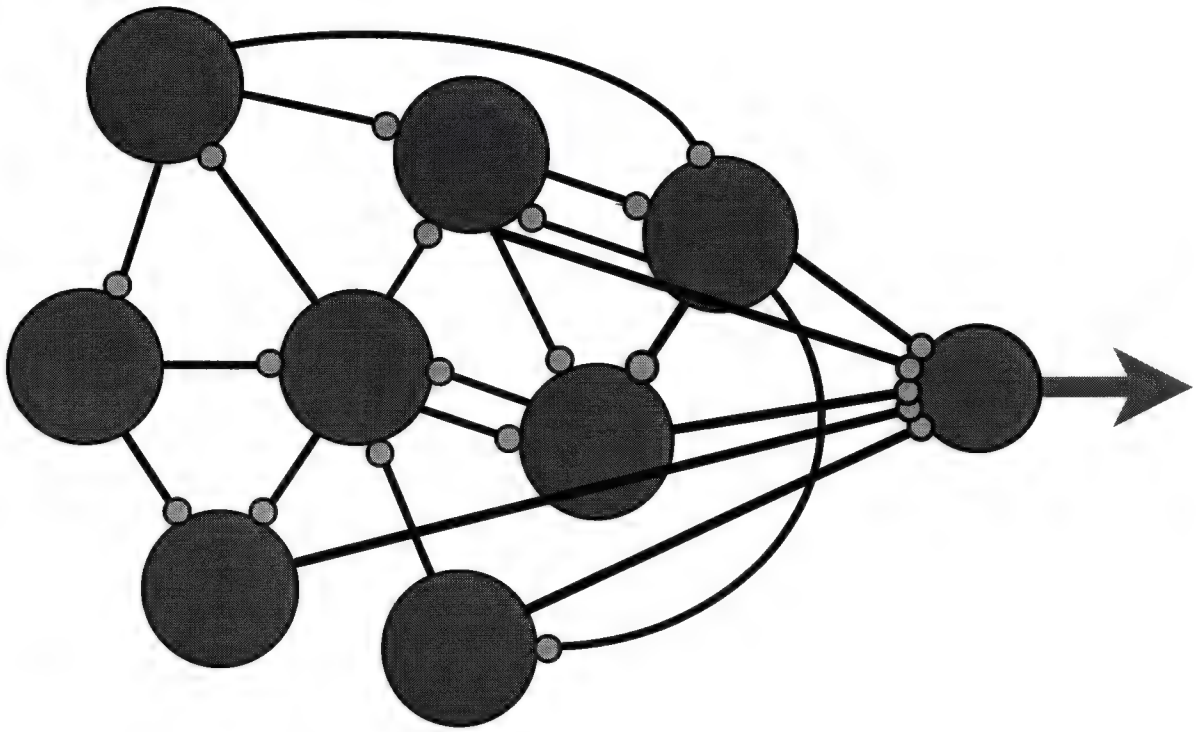

Maass, Jaeger

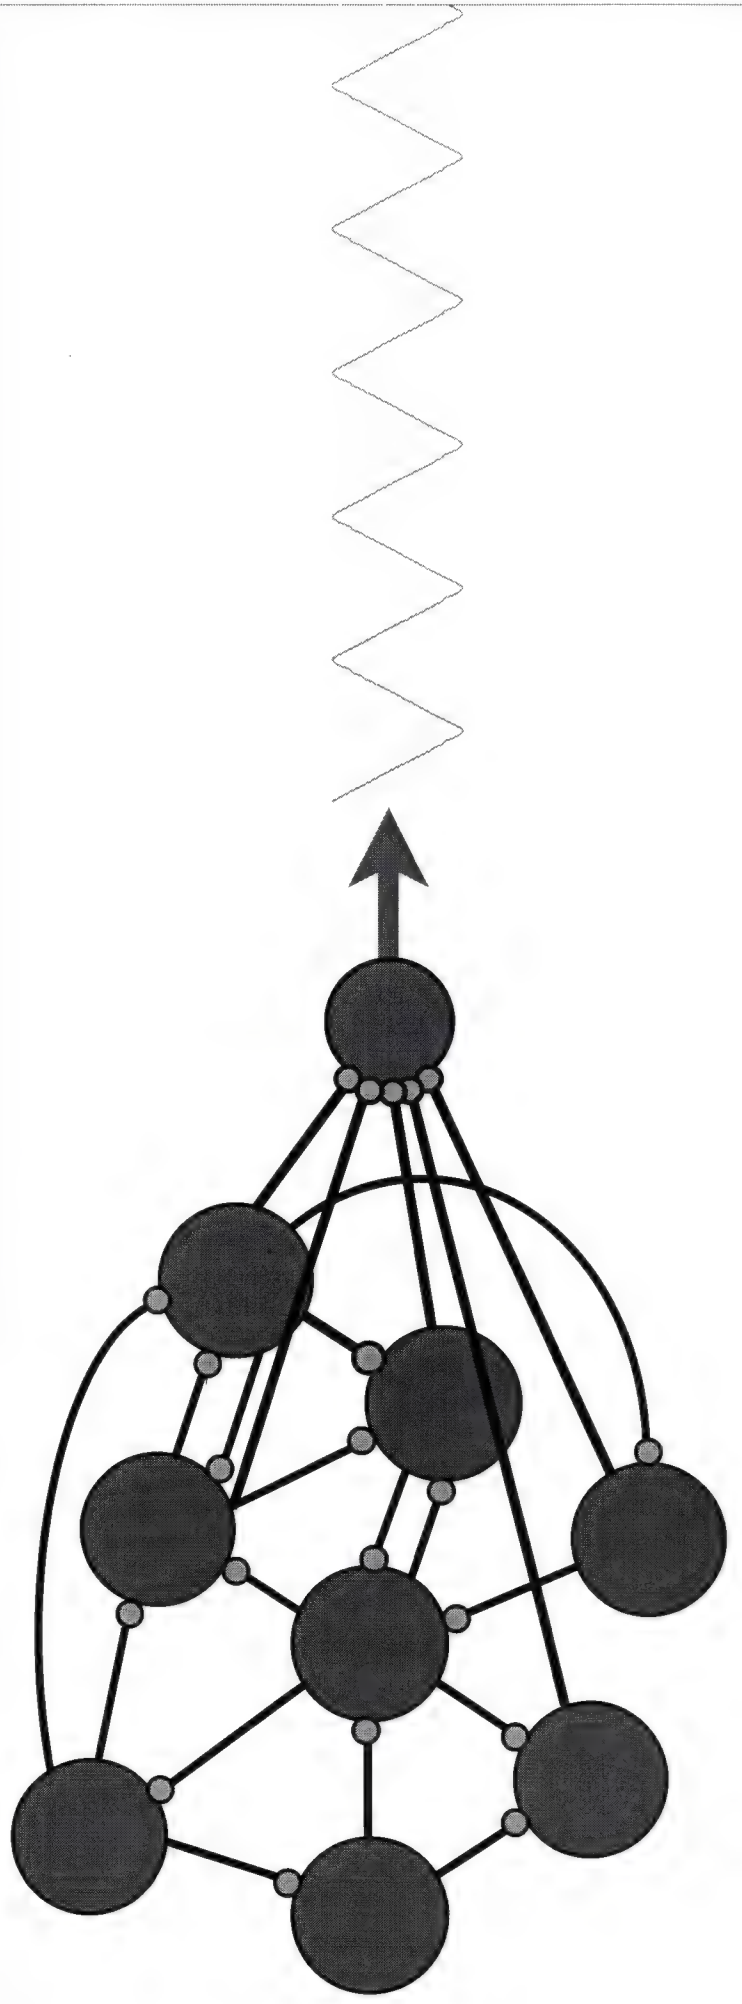

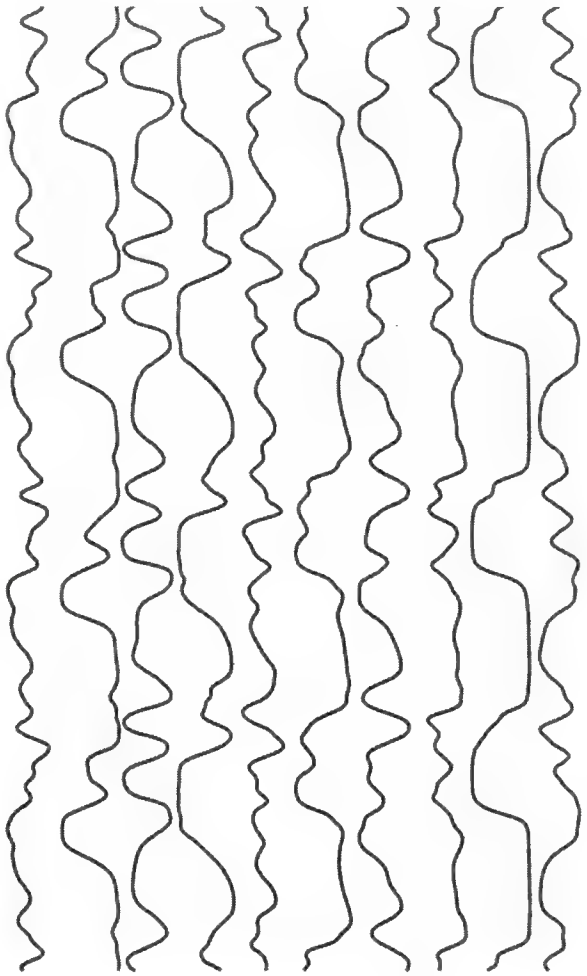

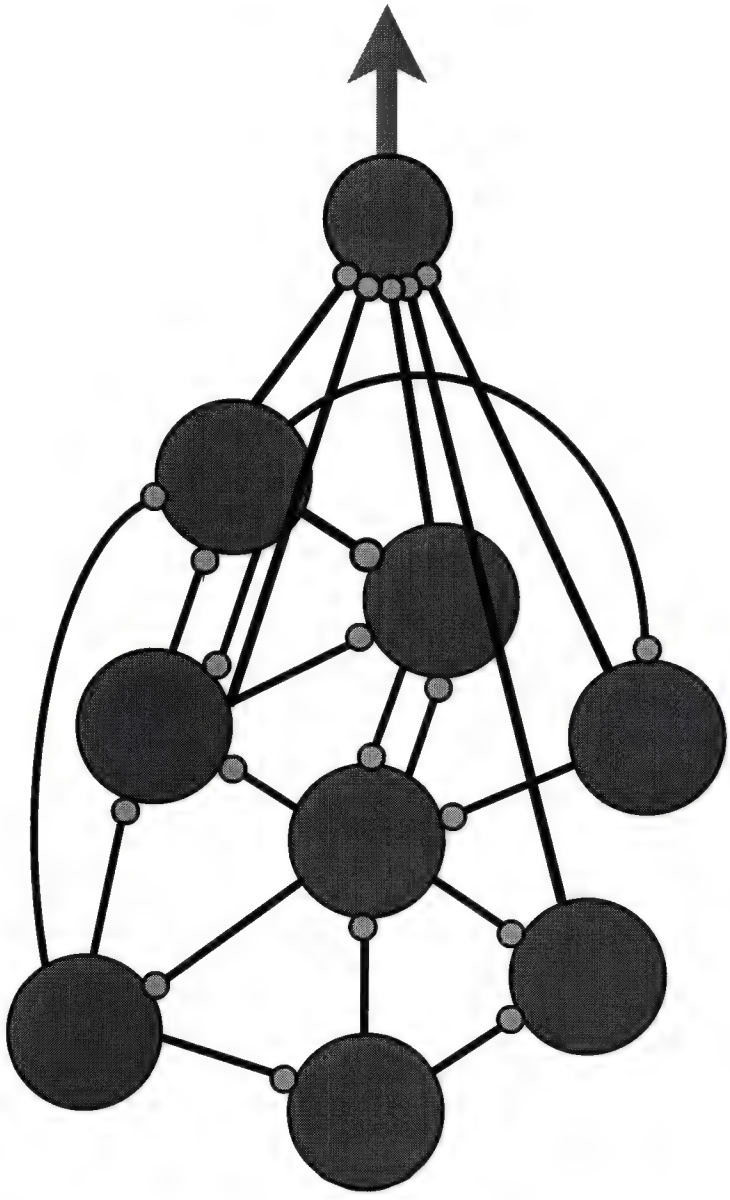

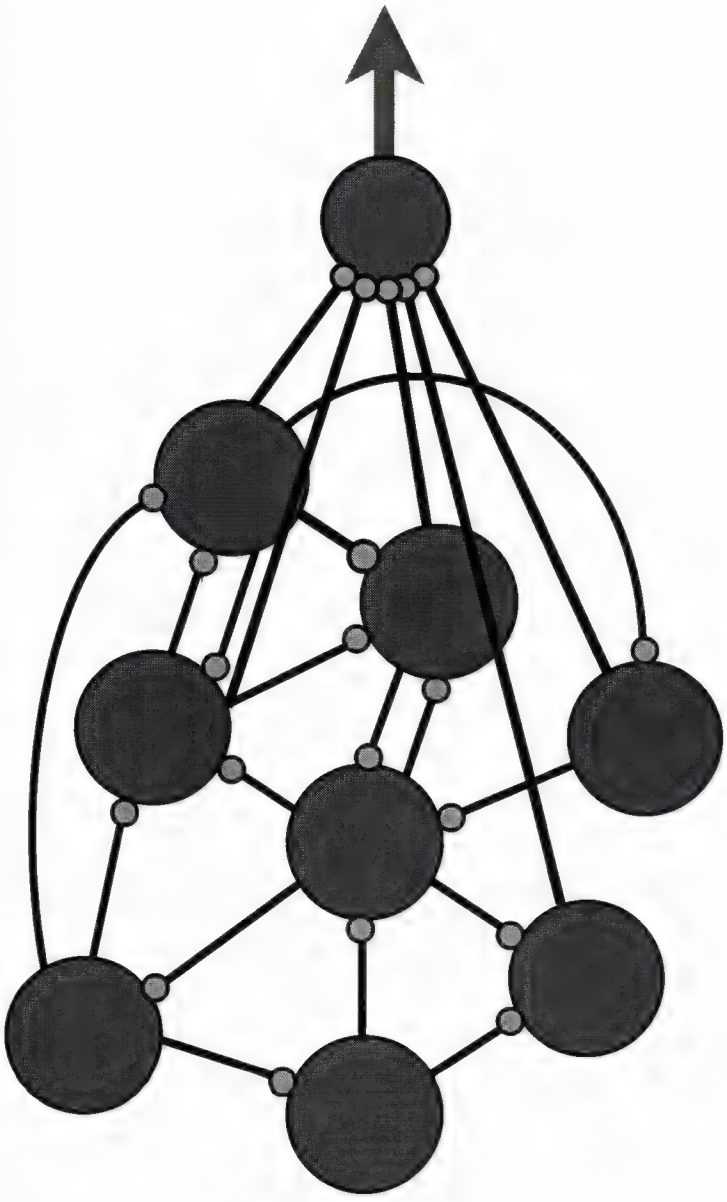

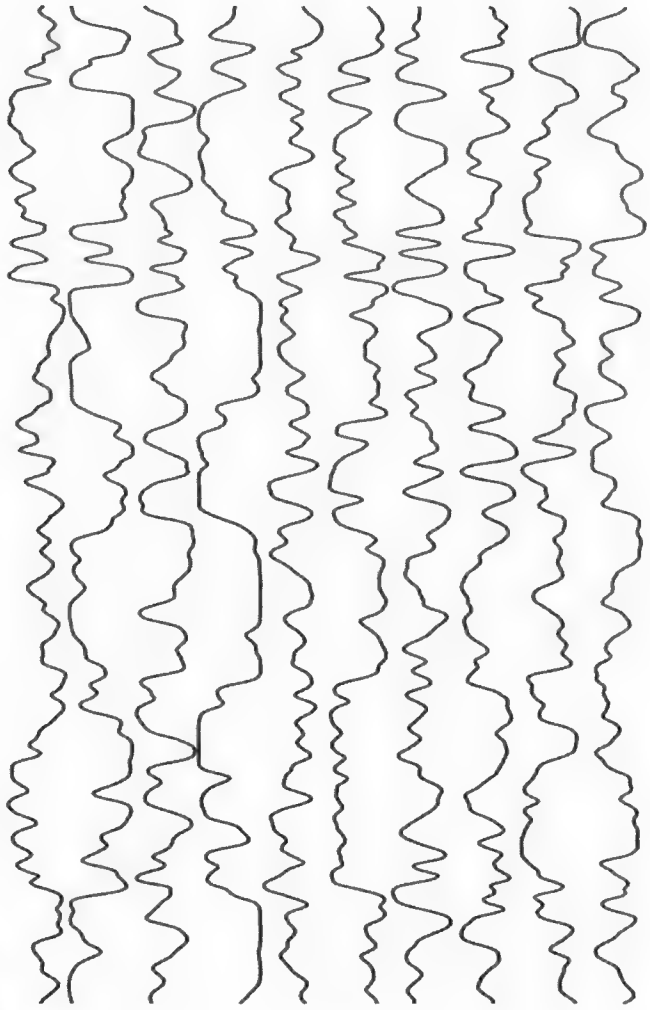

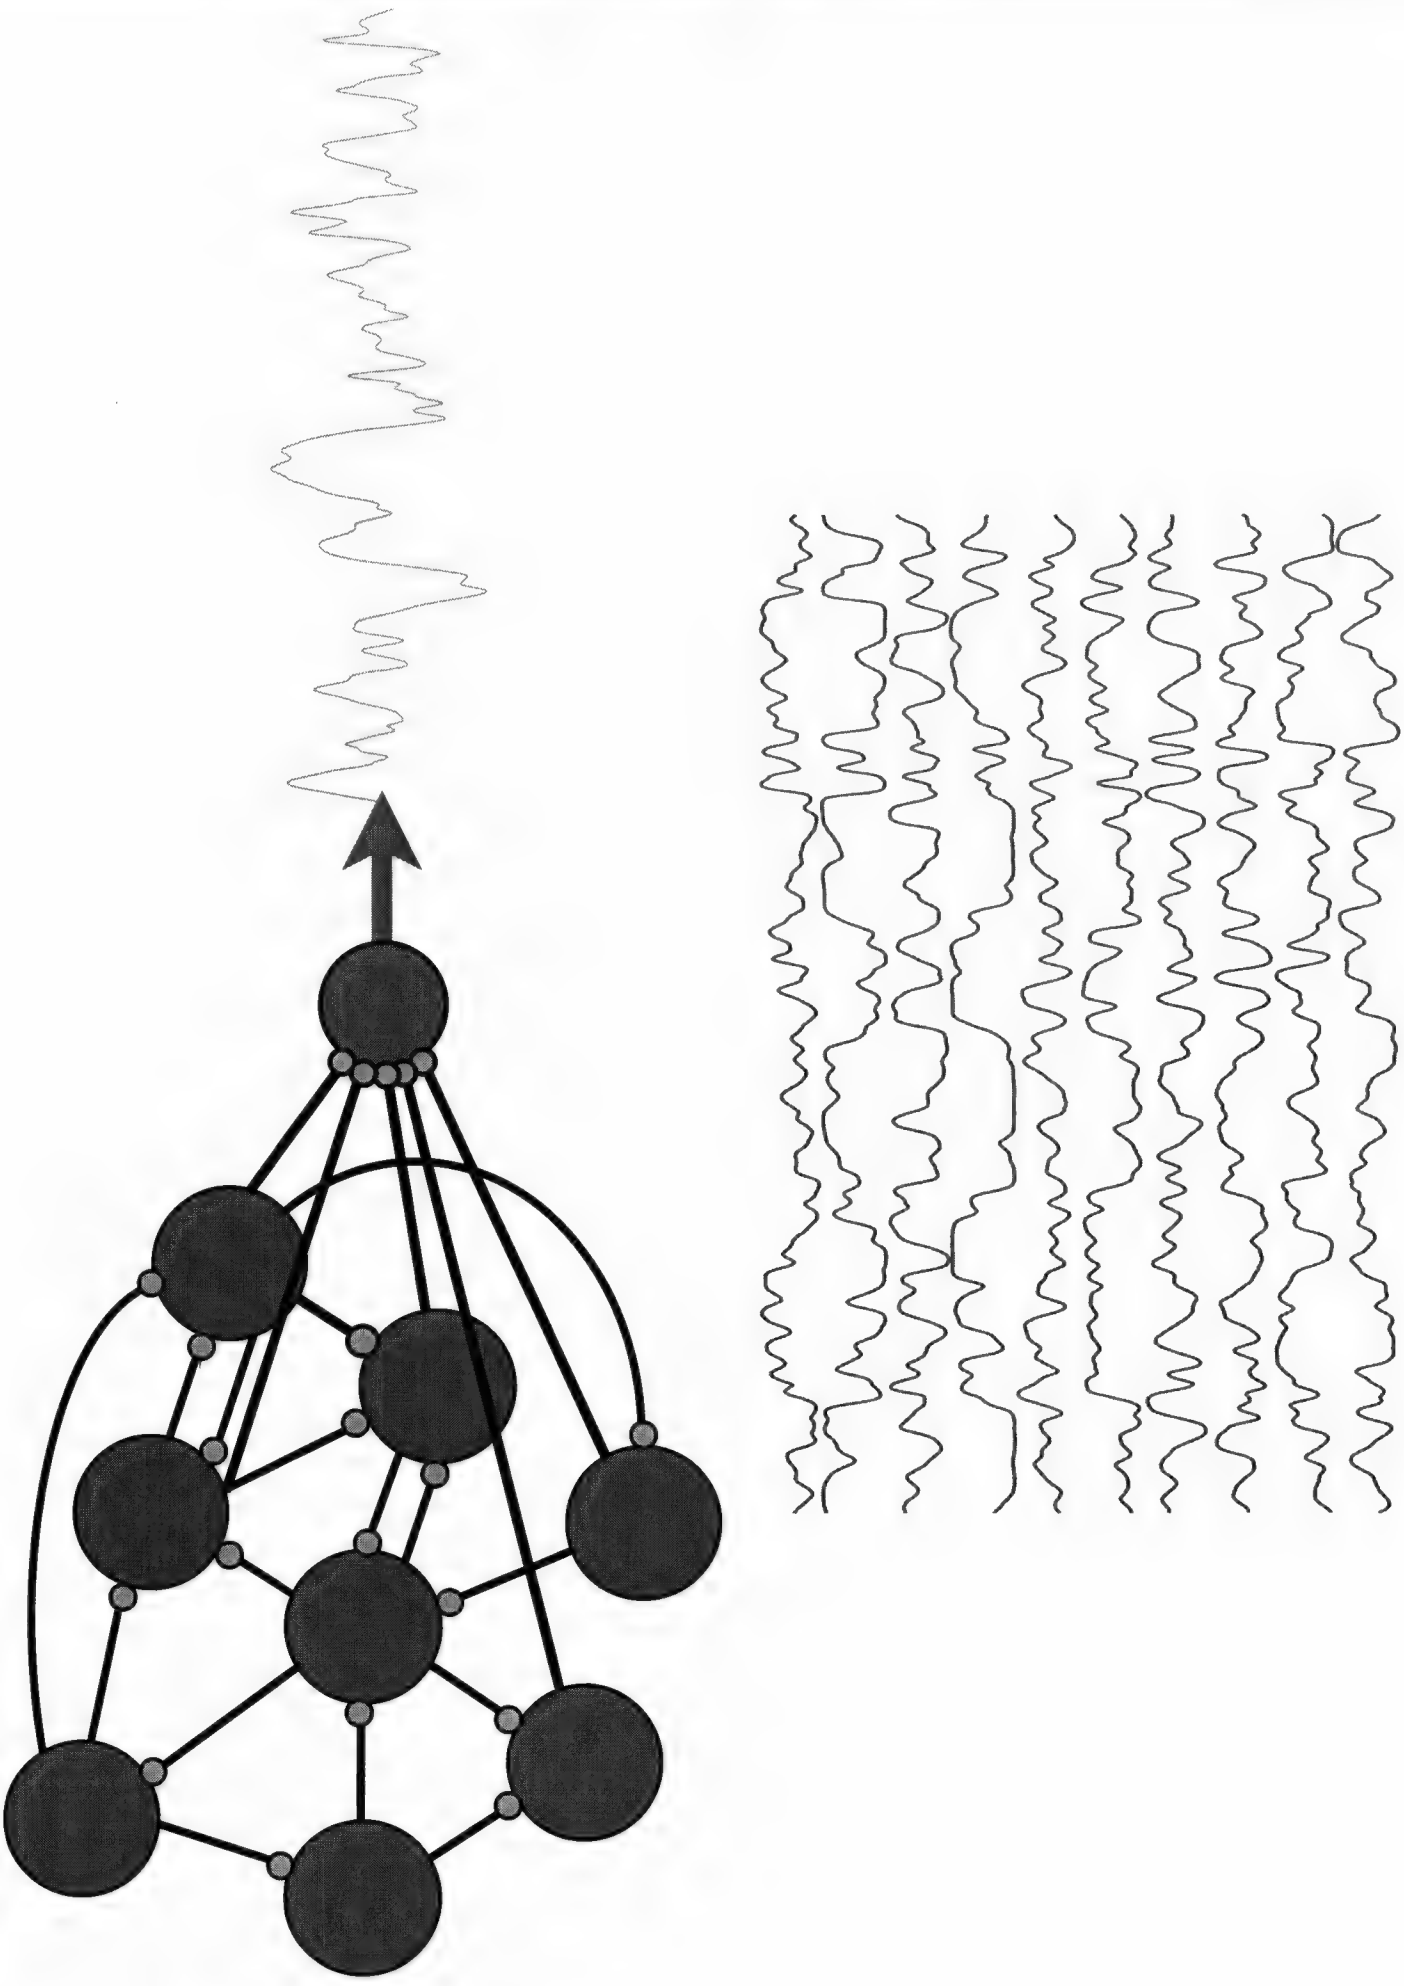

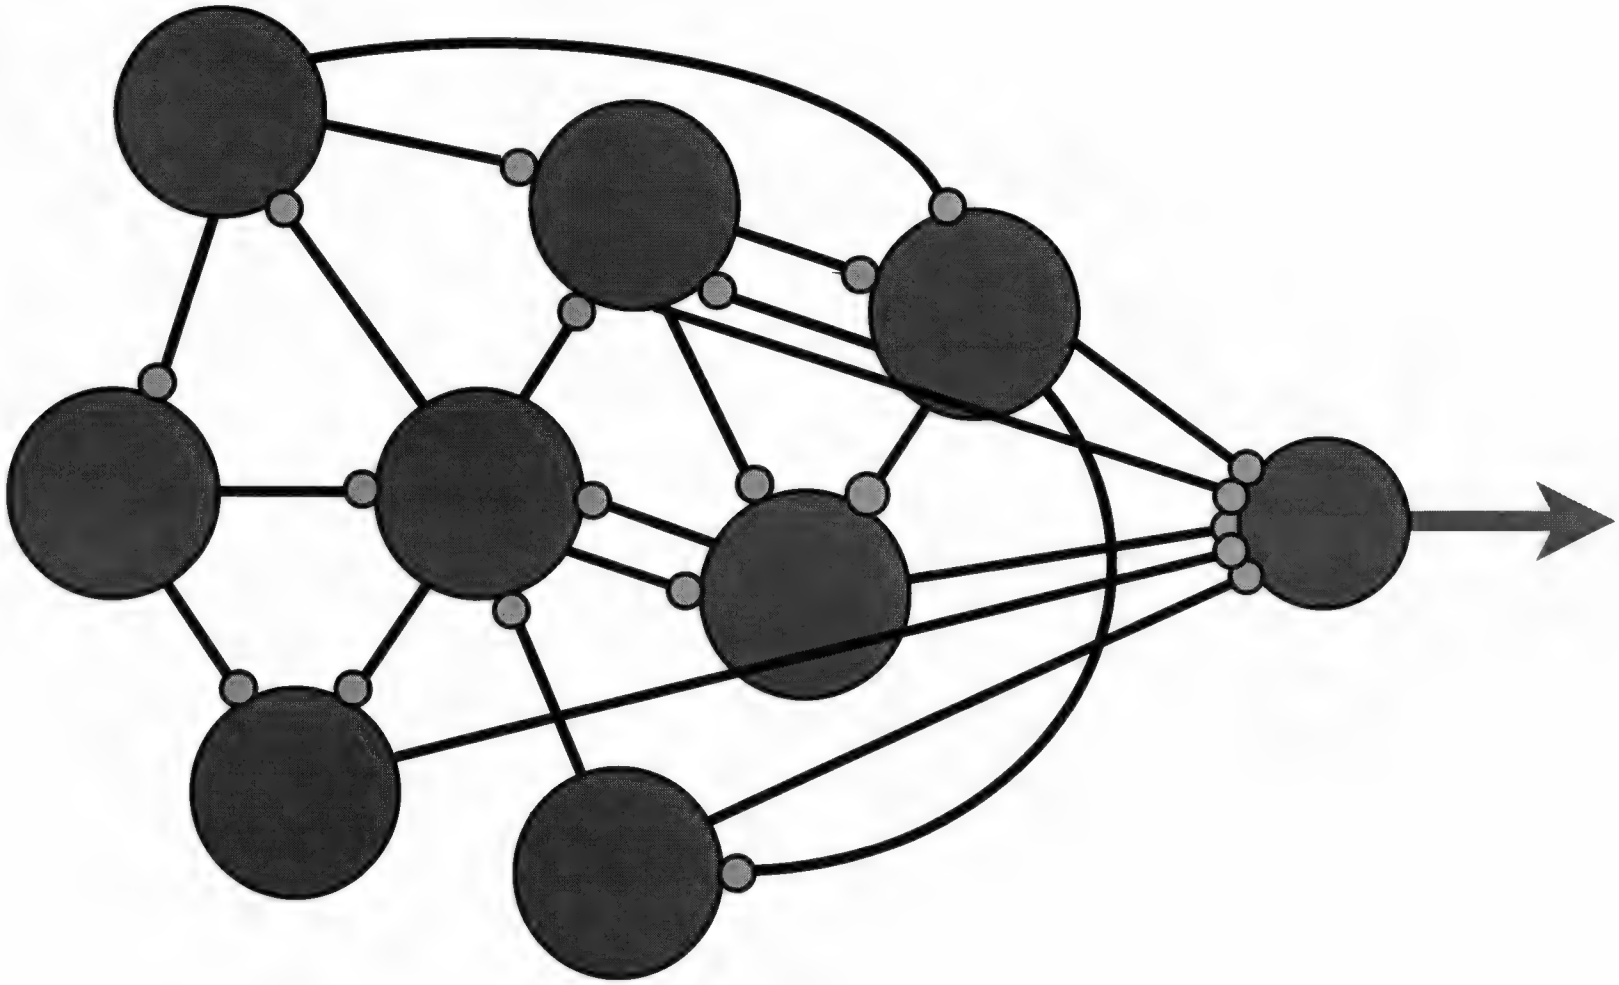

Jaeger and Haass, 2004

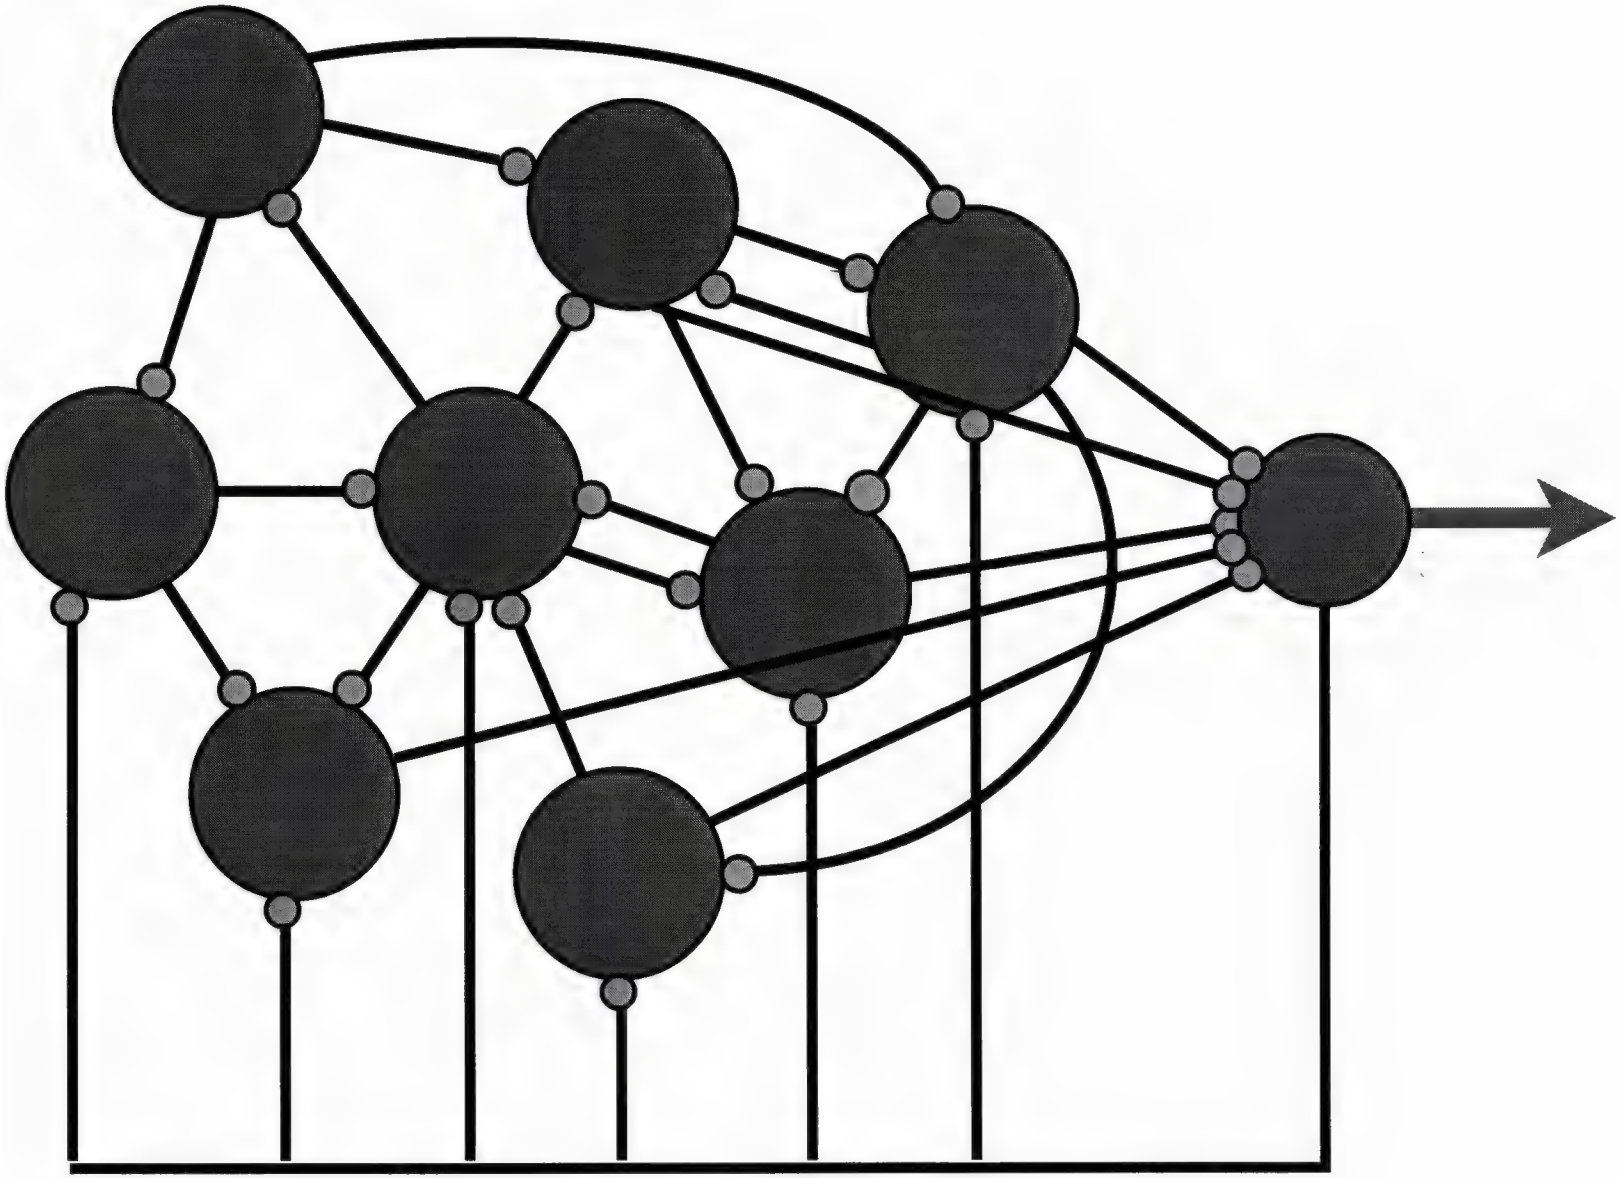

Jaeger and Haass, 2004

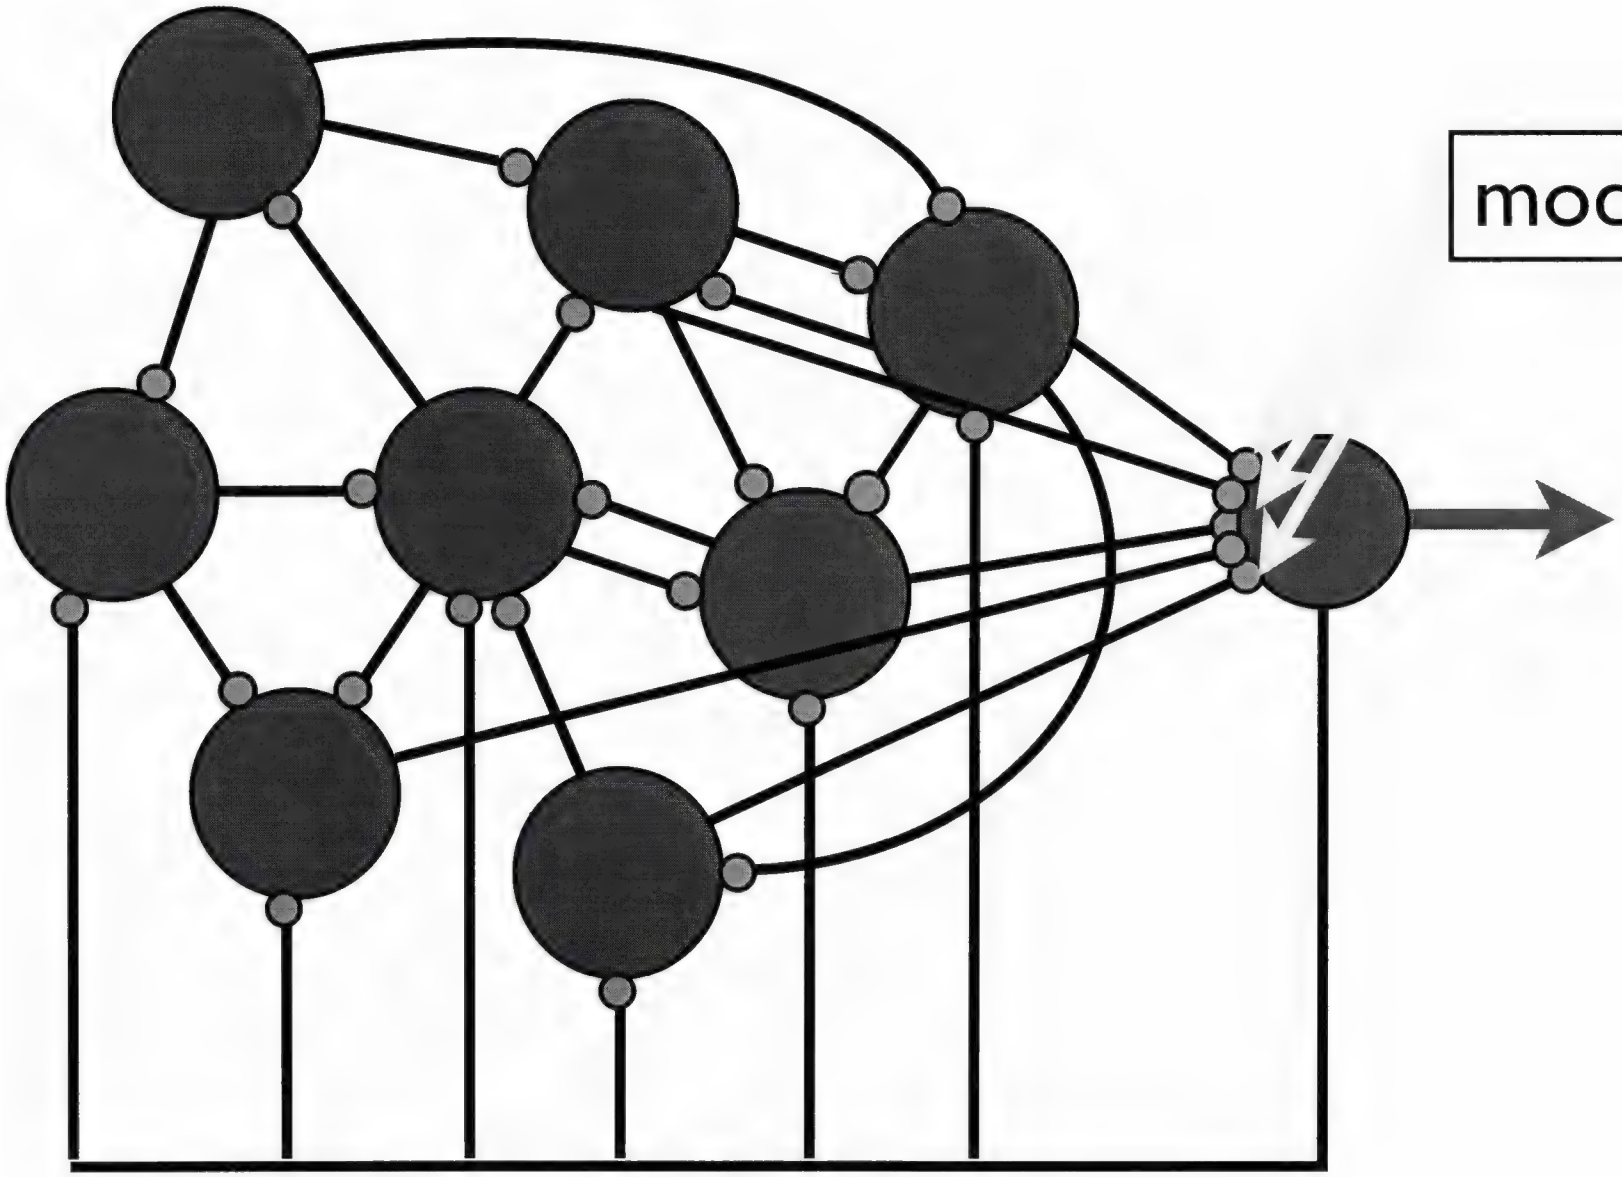

Jaeger and Haass, 2004

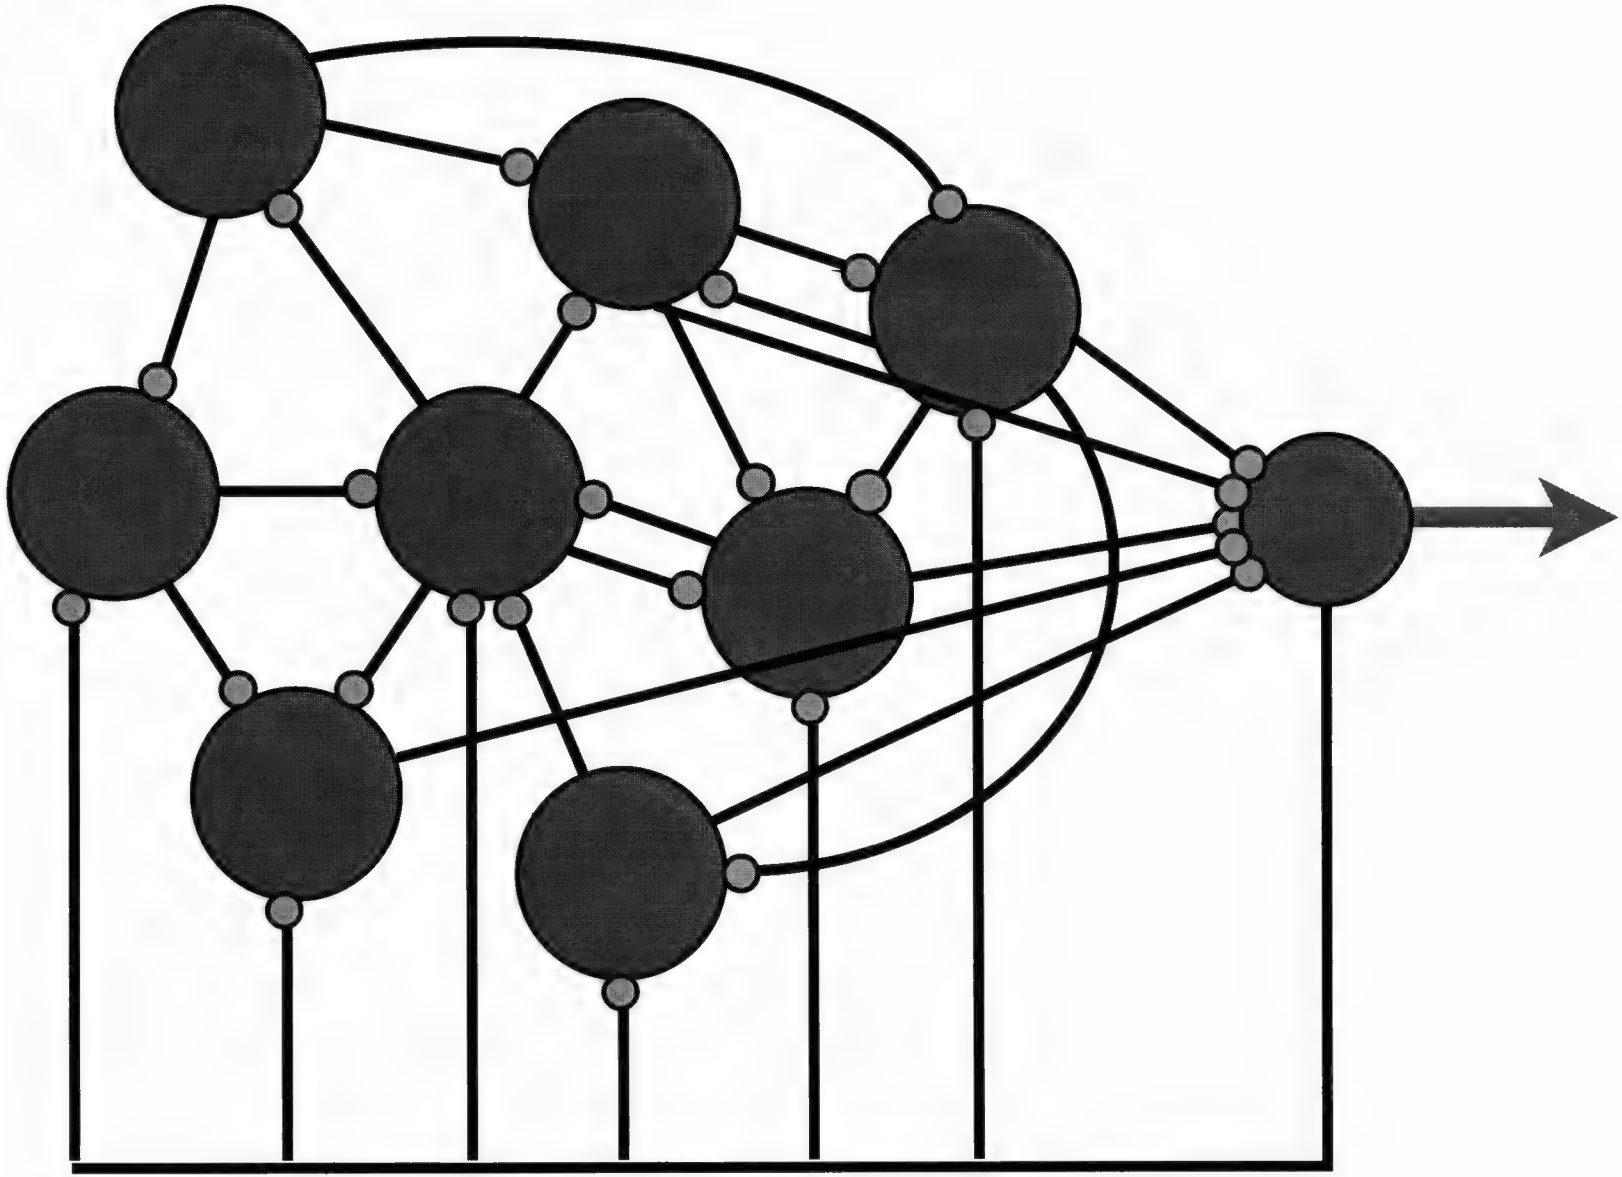

Jaeger and Haass, 2004

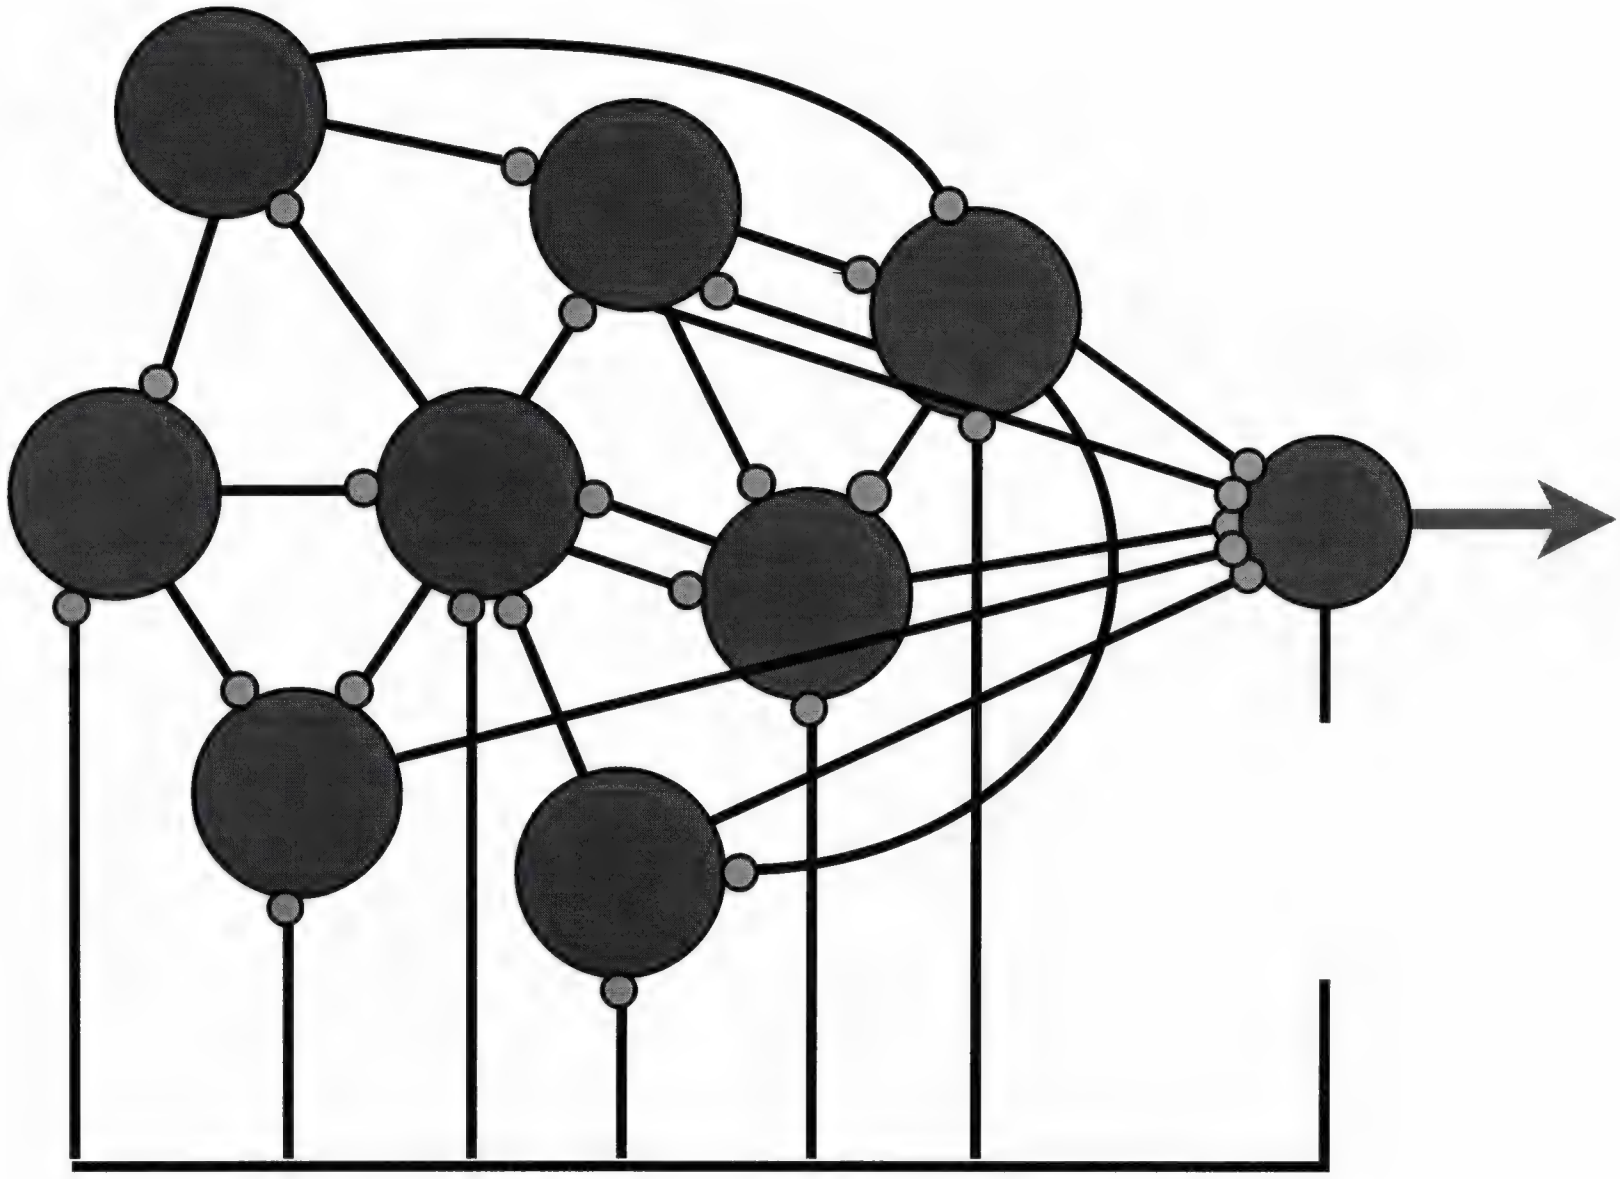

Jaeger and Haass, 2004

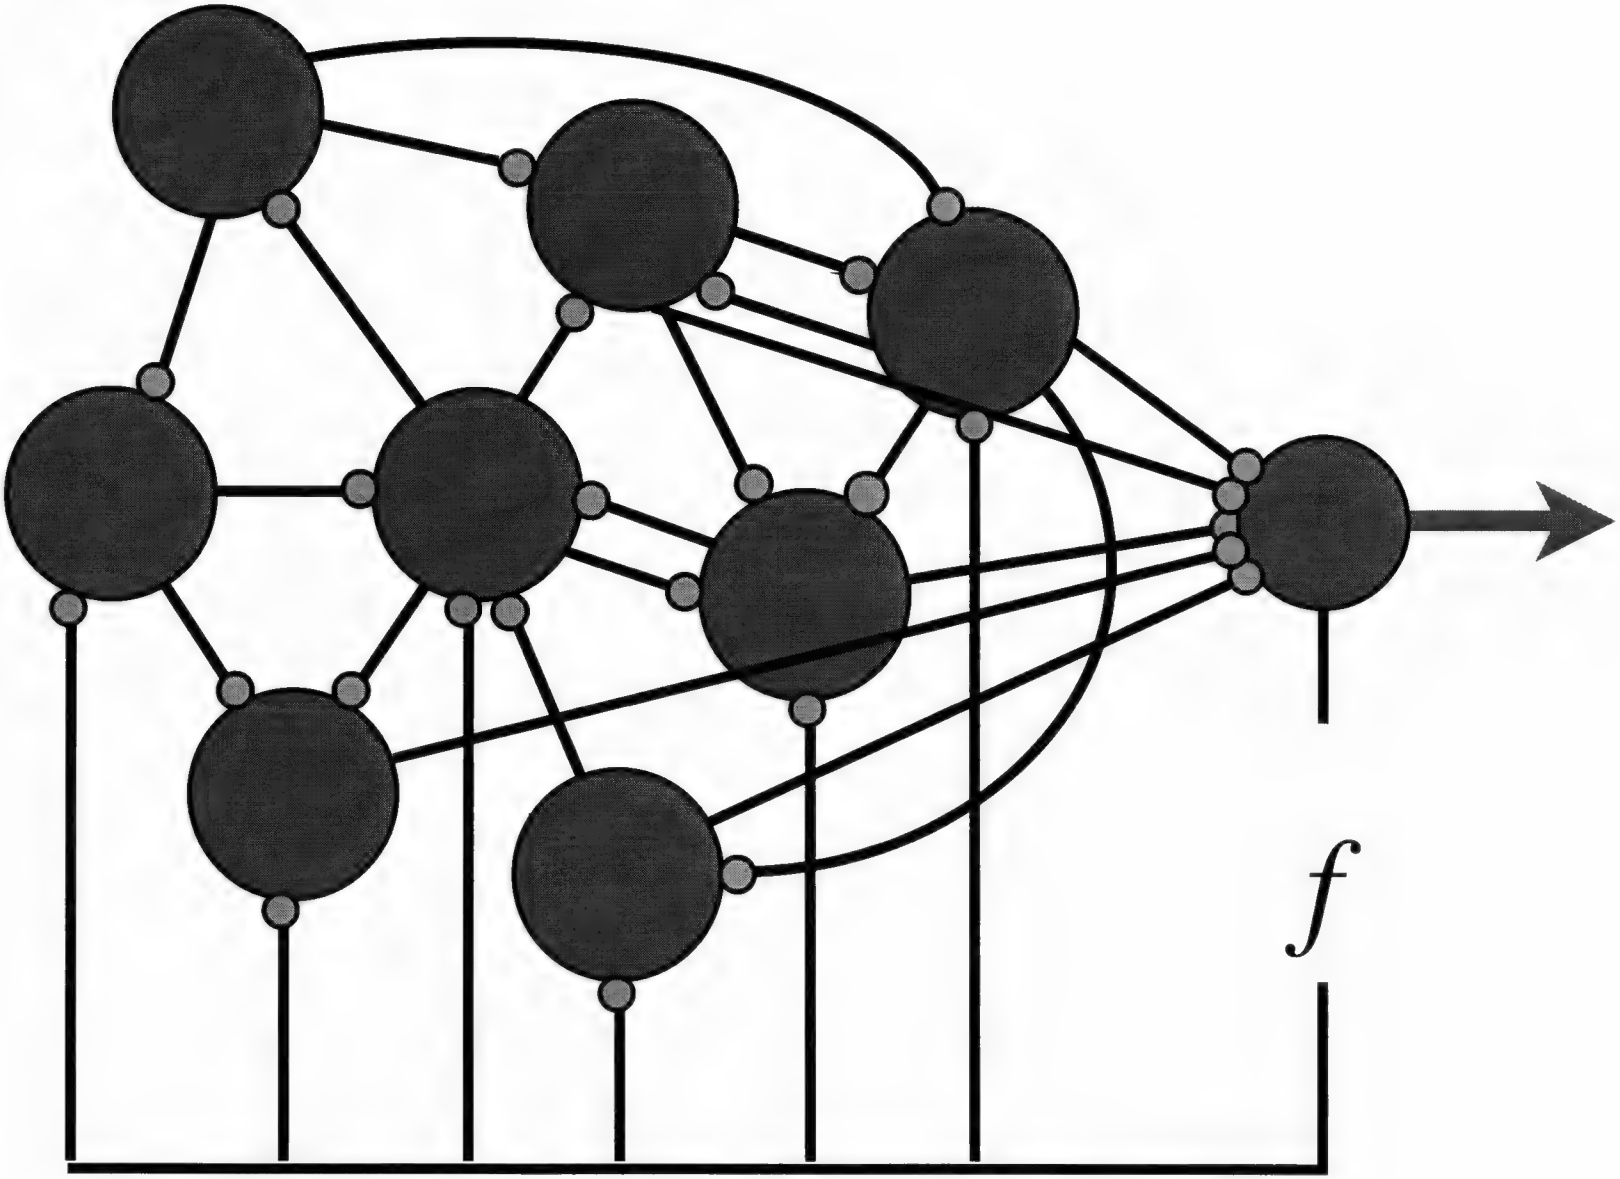

Jaeger and Haass, 2004

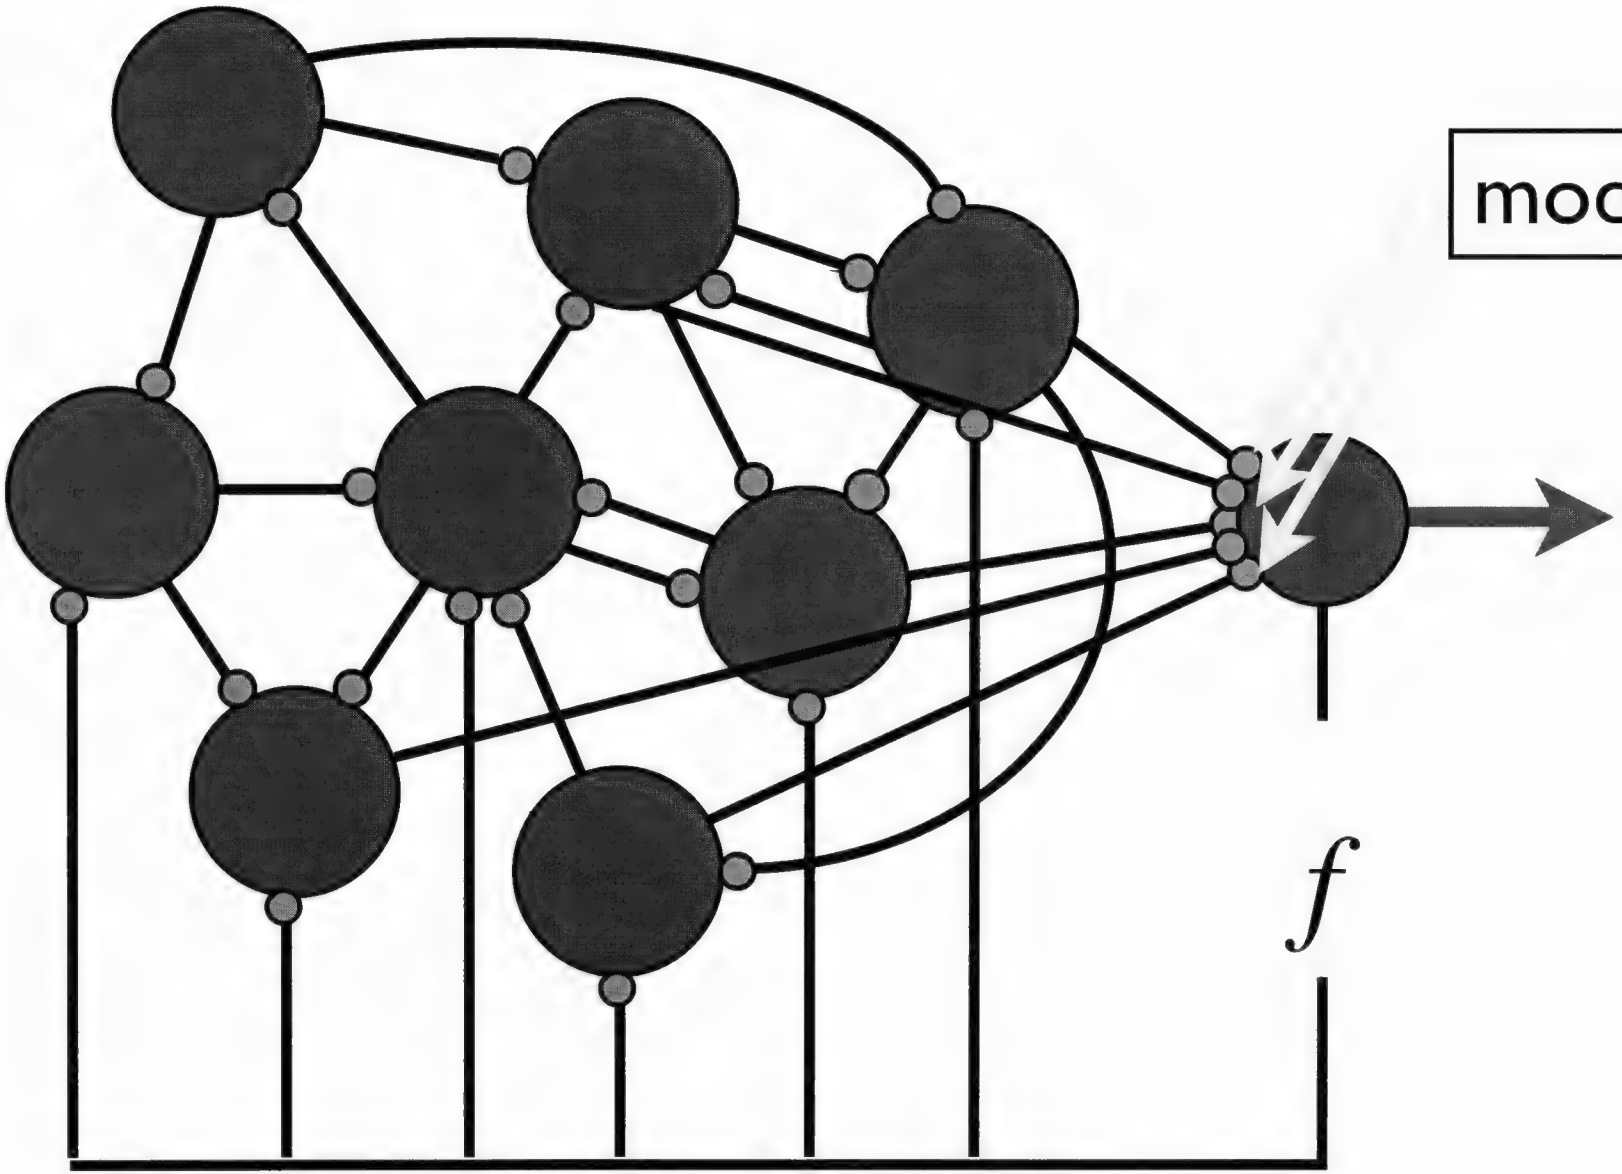

Jaeger and Haass, 2004

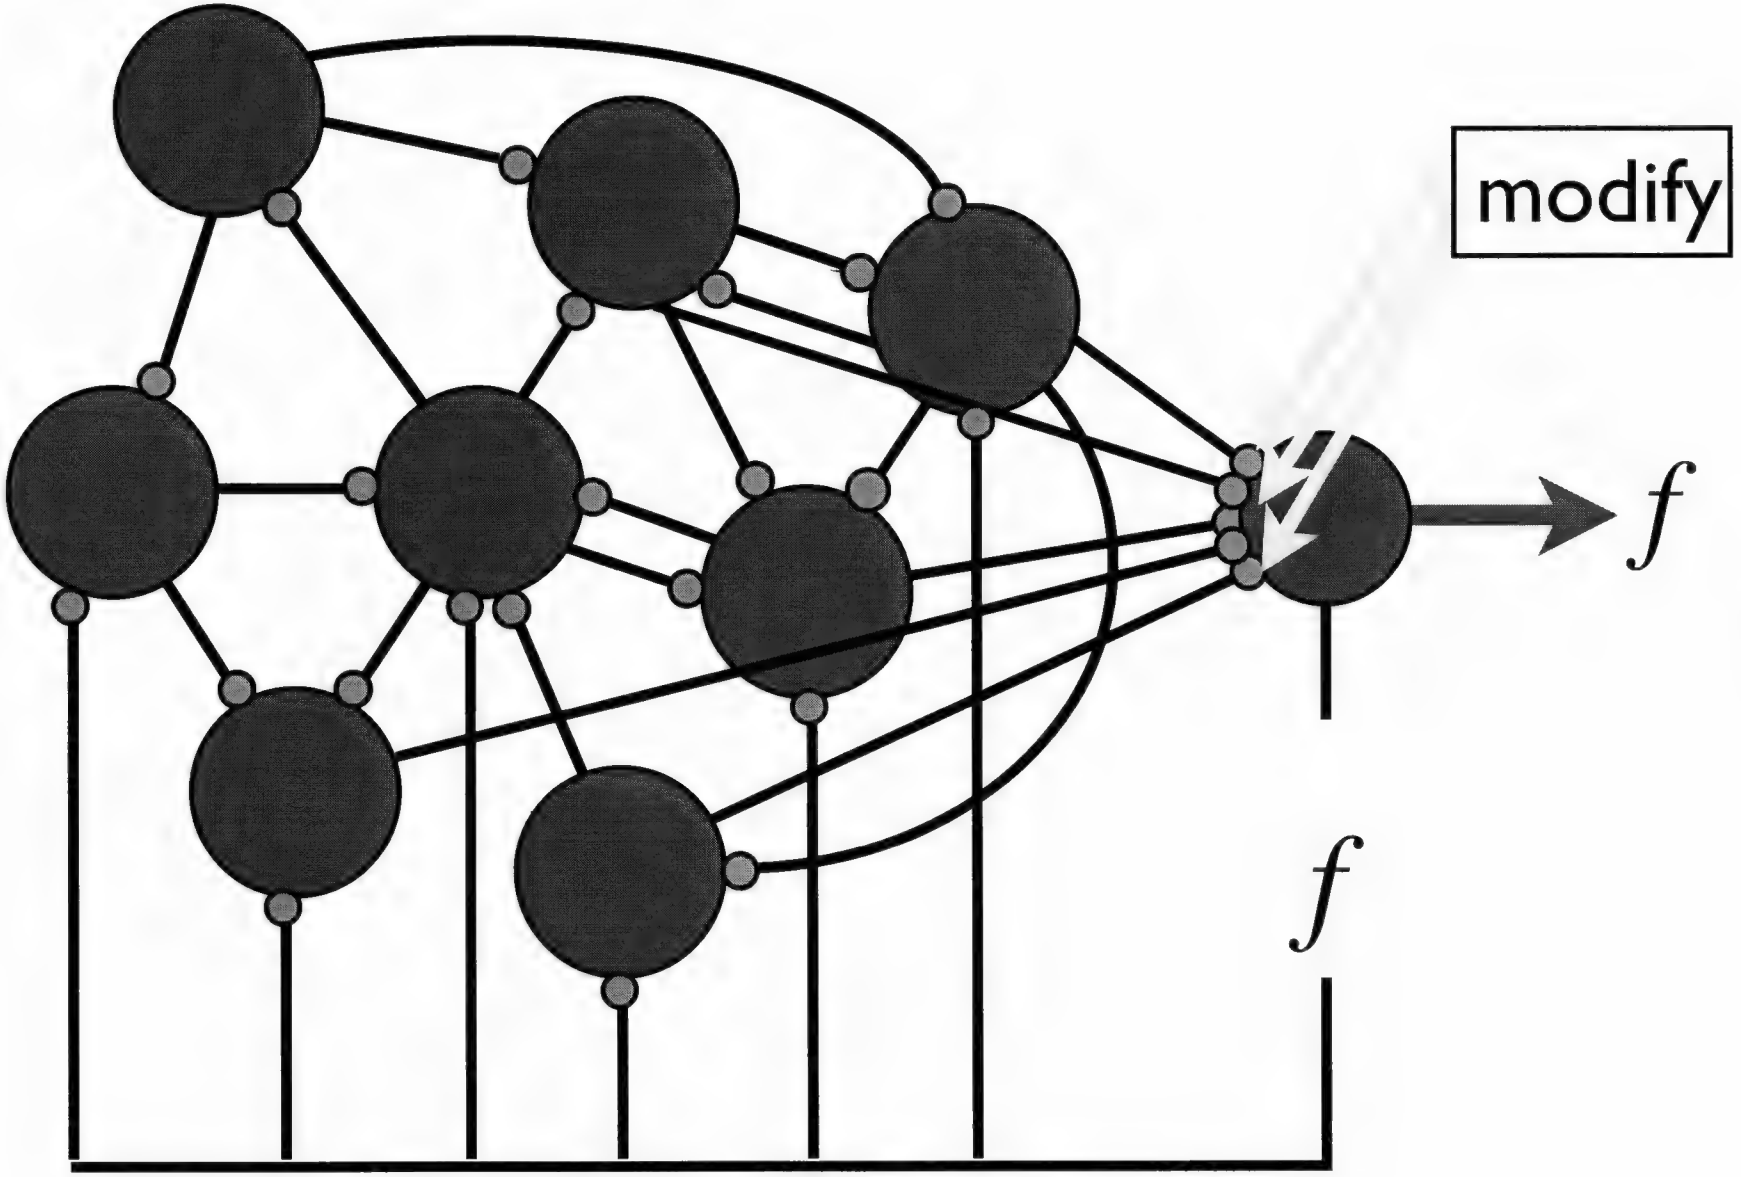

Jaeger and Haass, 2004

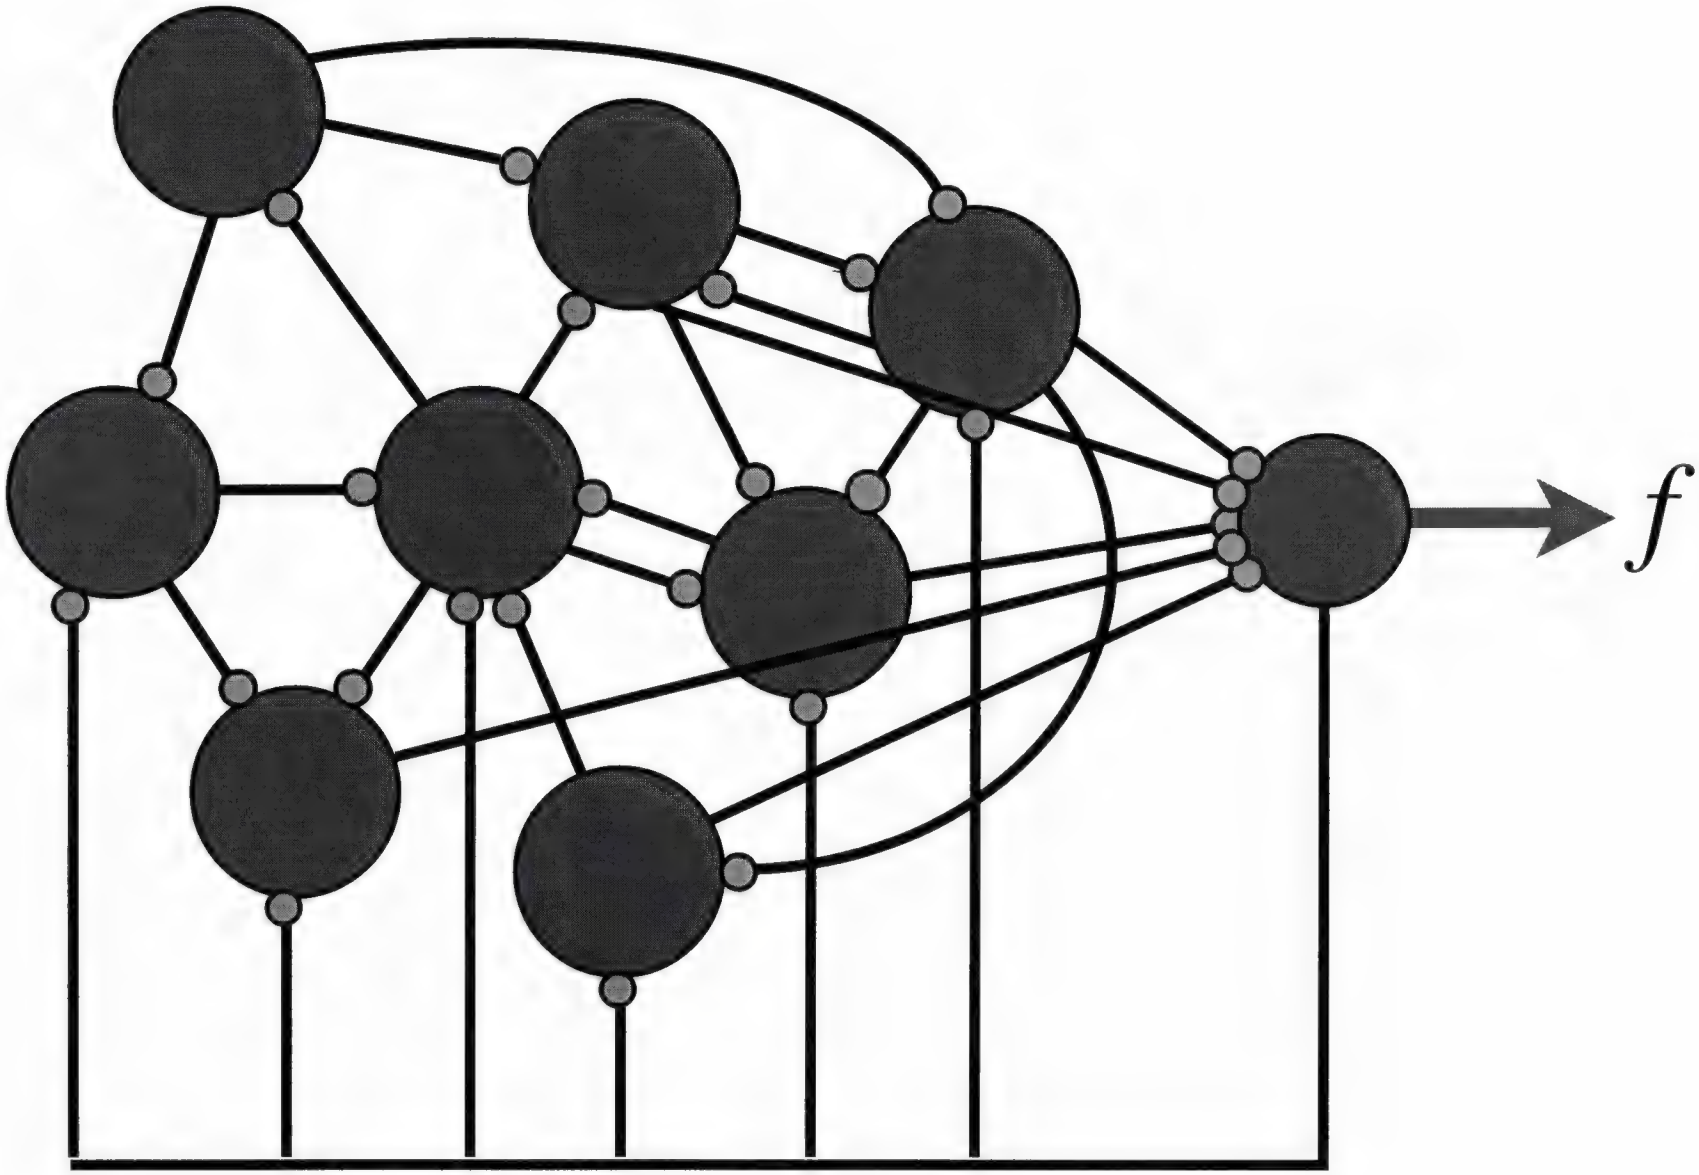

Jaeger and Haass, 2004

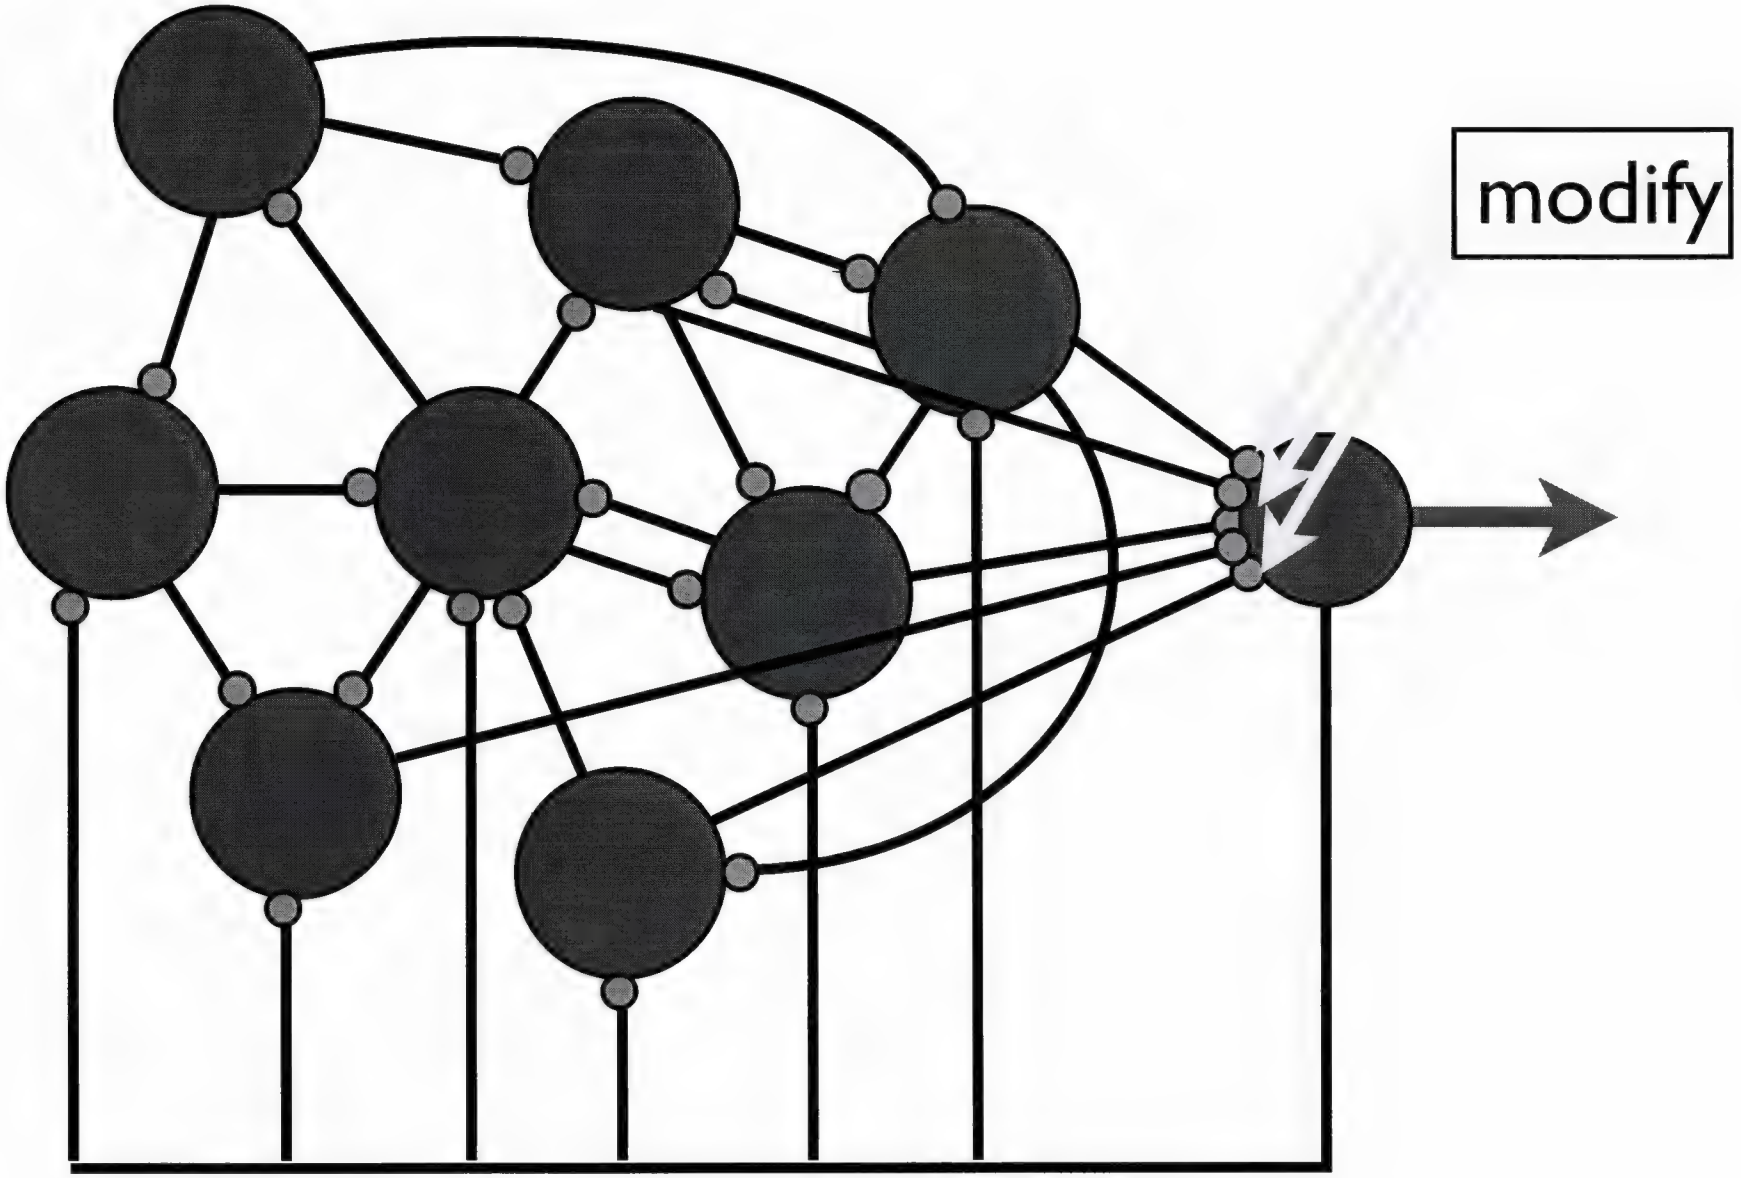
modify

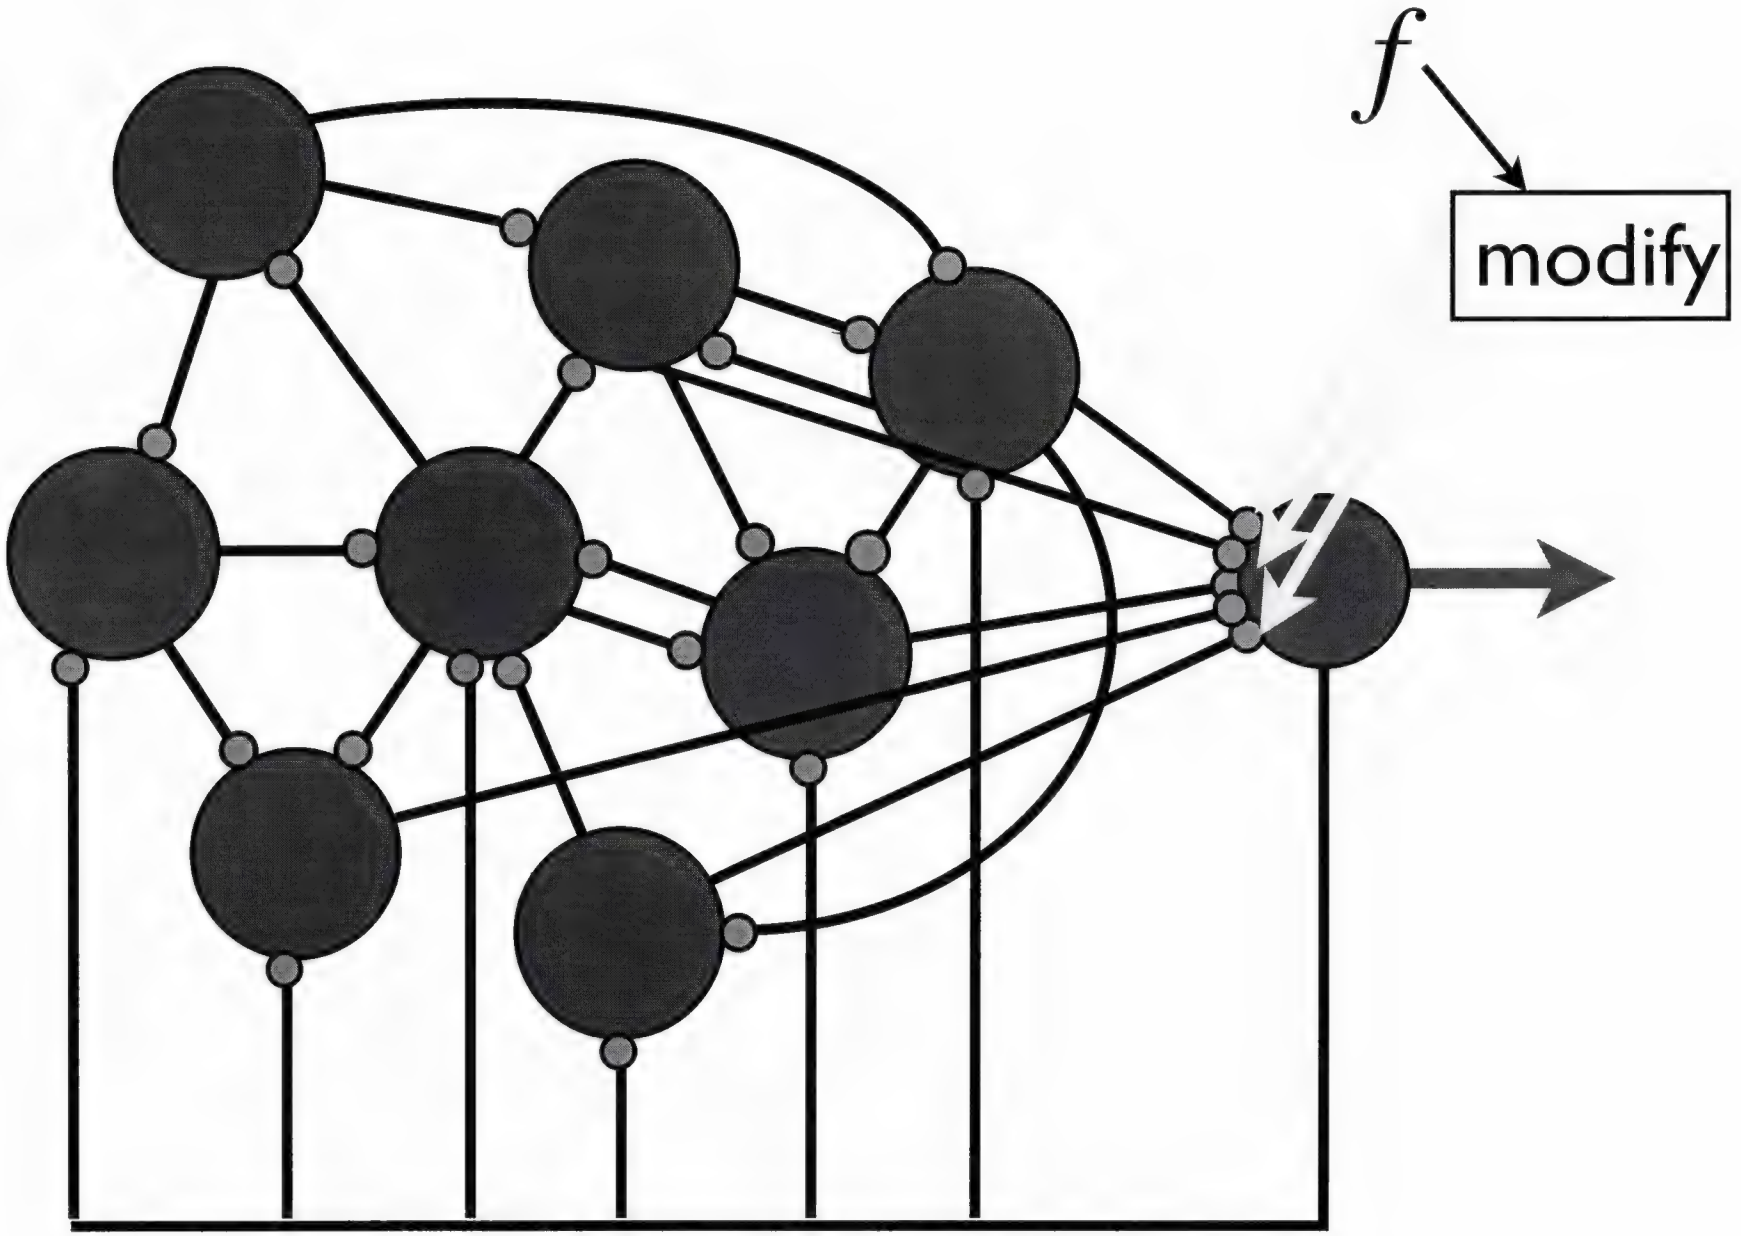
f
modify

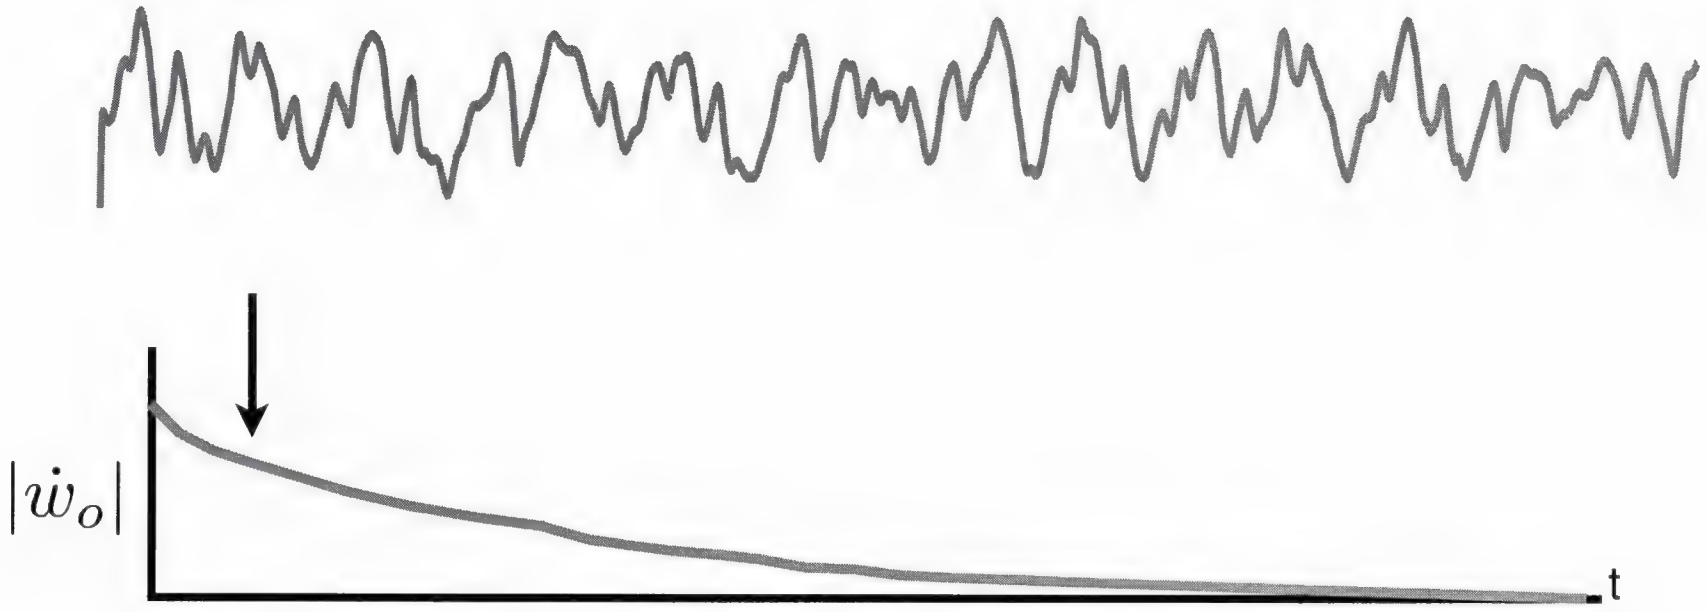
$|\dot{w}_o|$
t

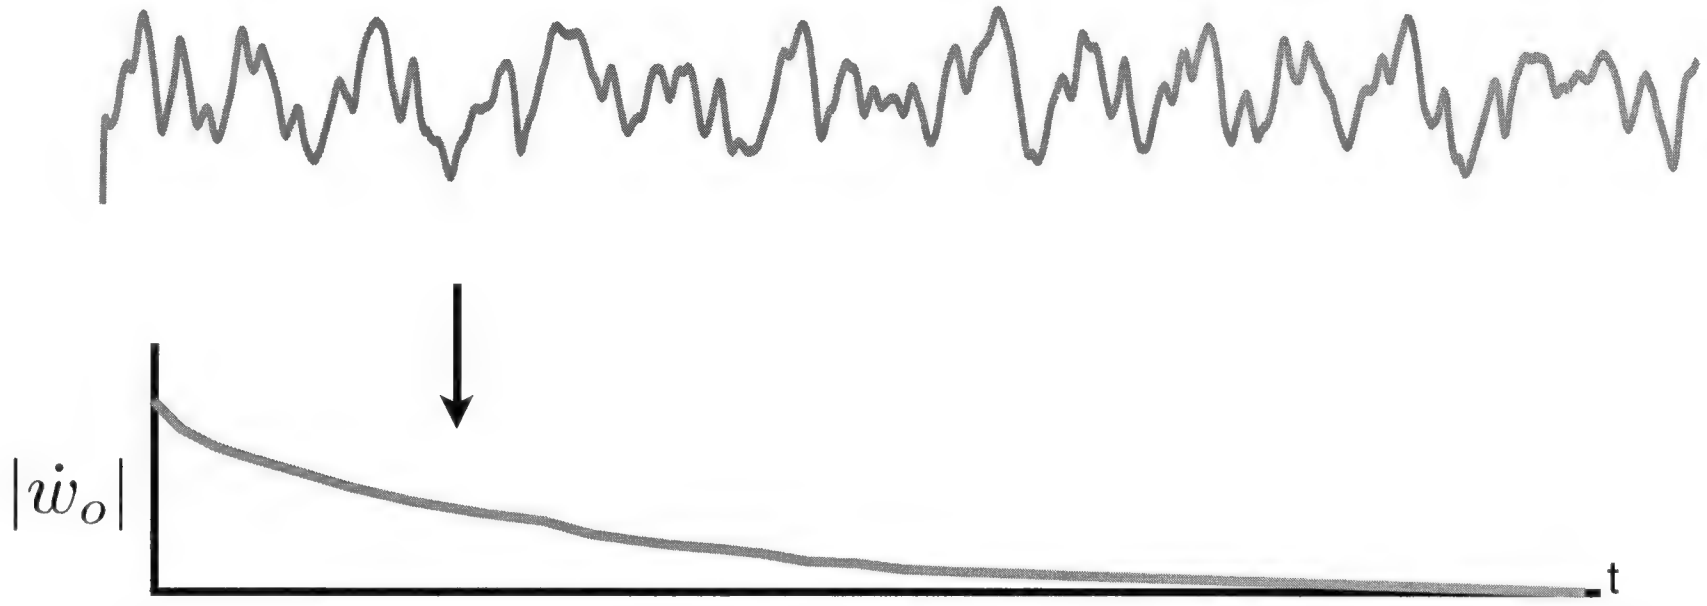
$|\dot{w}_o|$
t

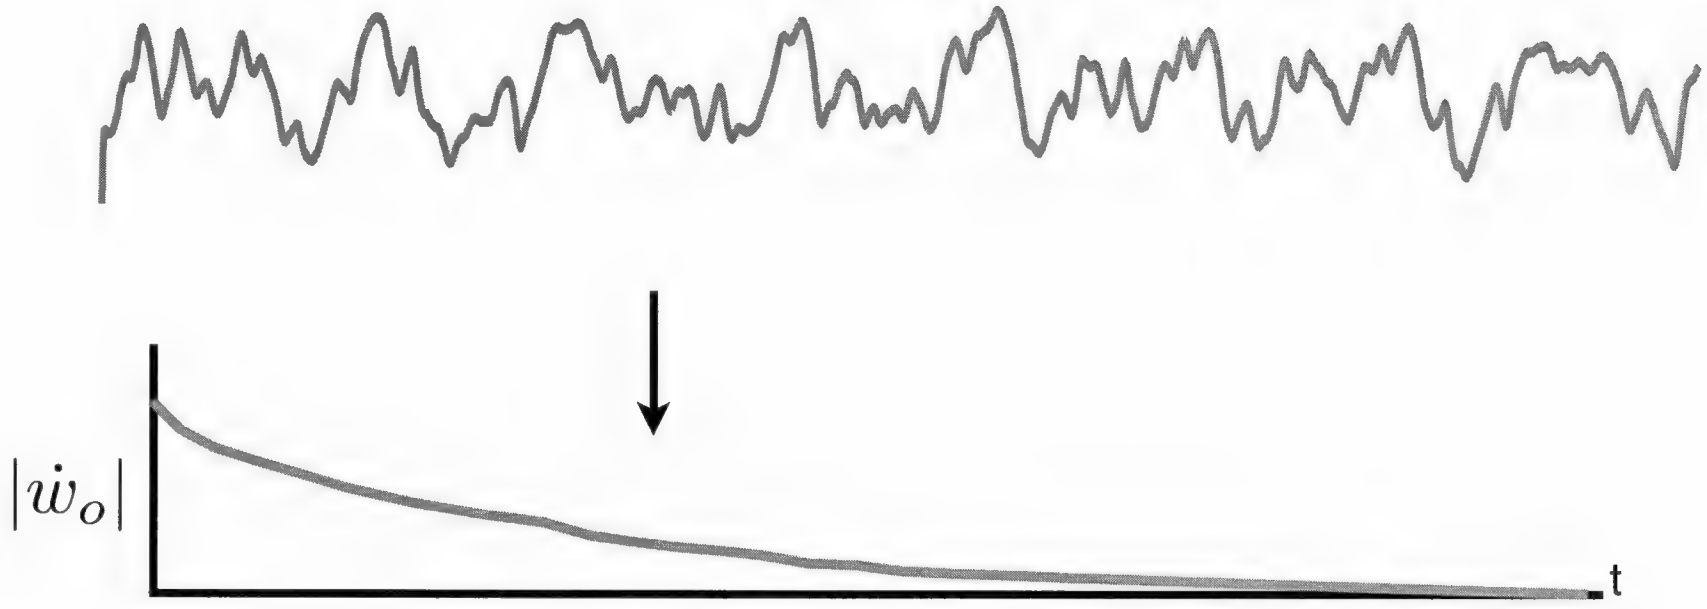
$|\dot{w}_o|$
t

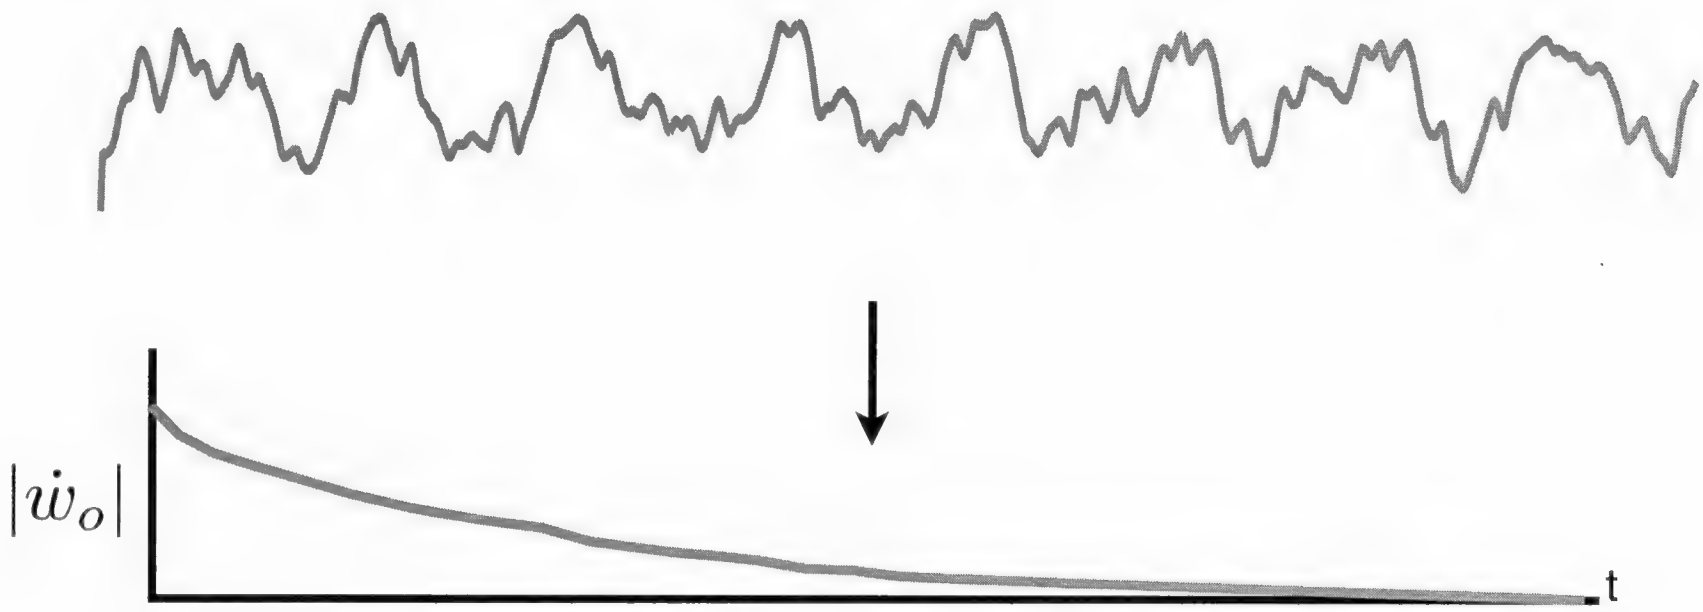
$|\dot{w}_o|$
t

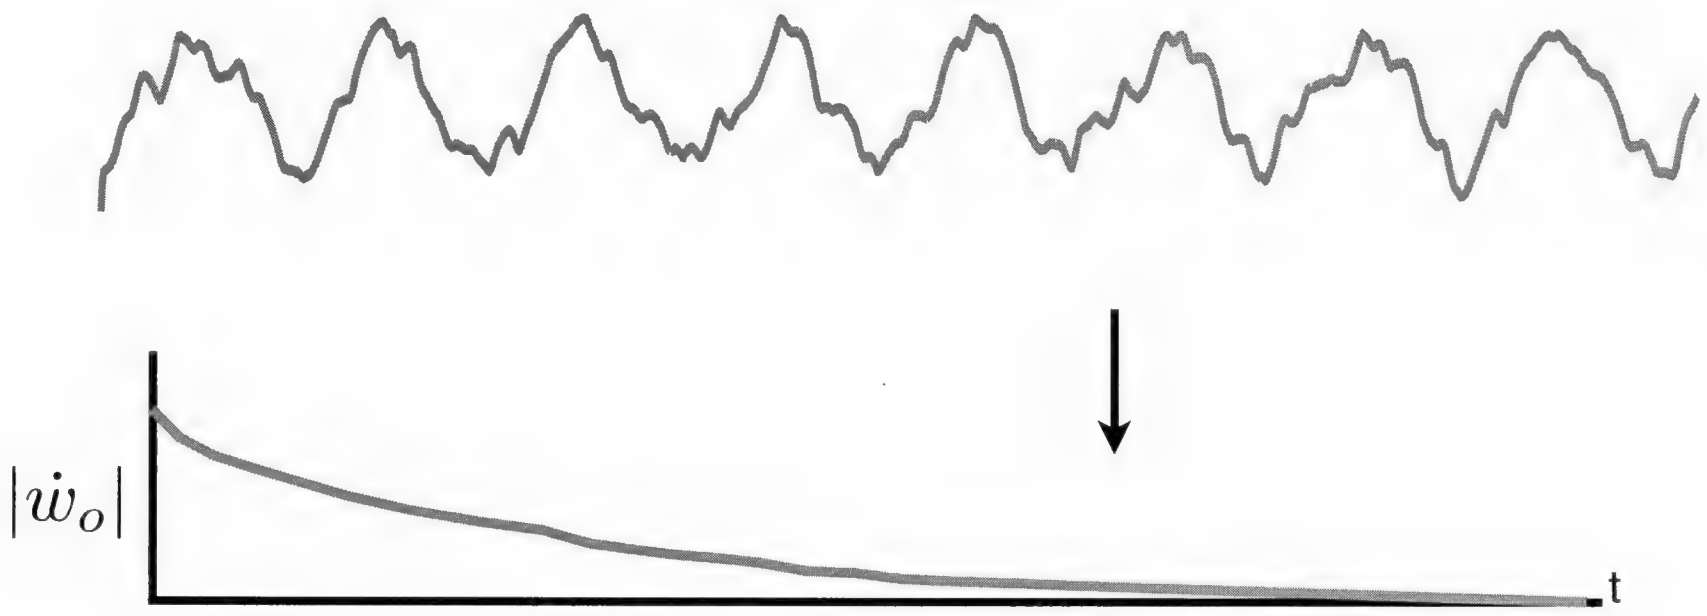
$|\dot{w}_o|$
t

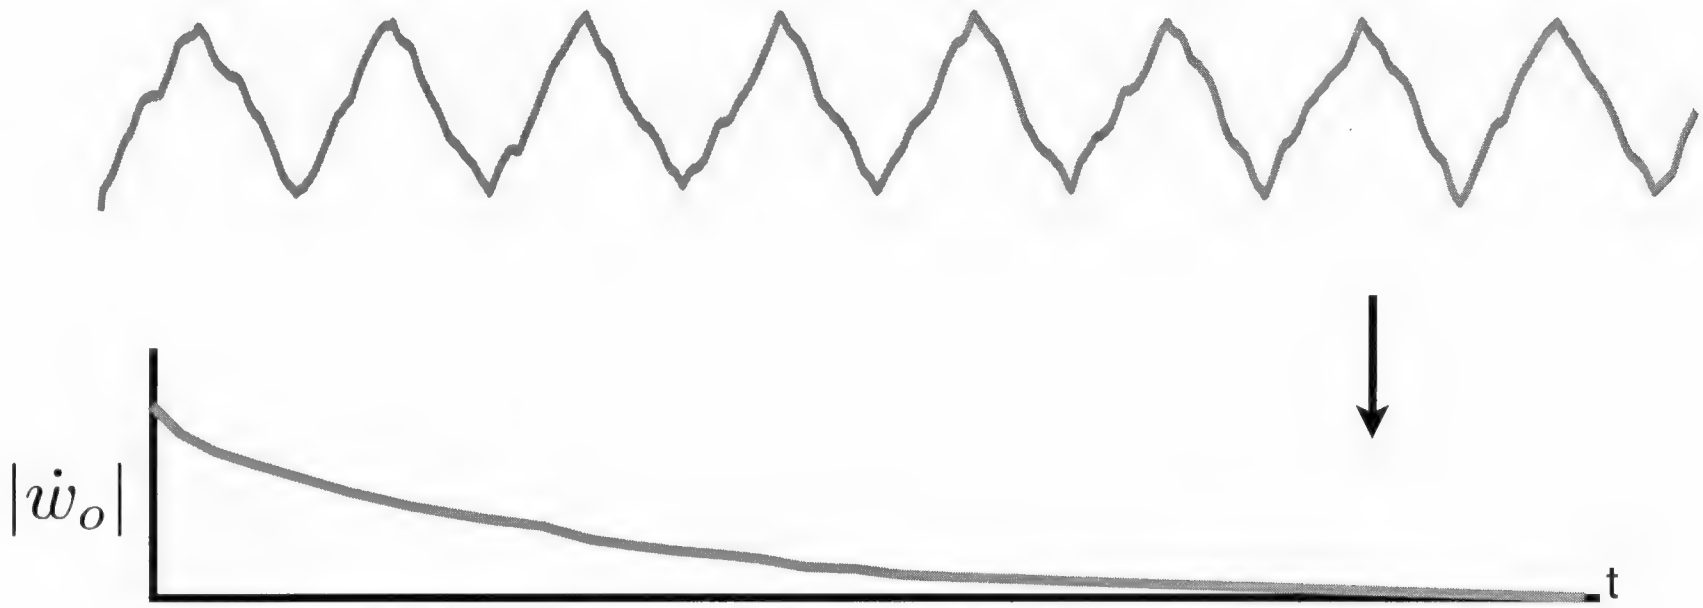
$|\dot{w}_o|$
t

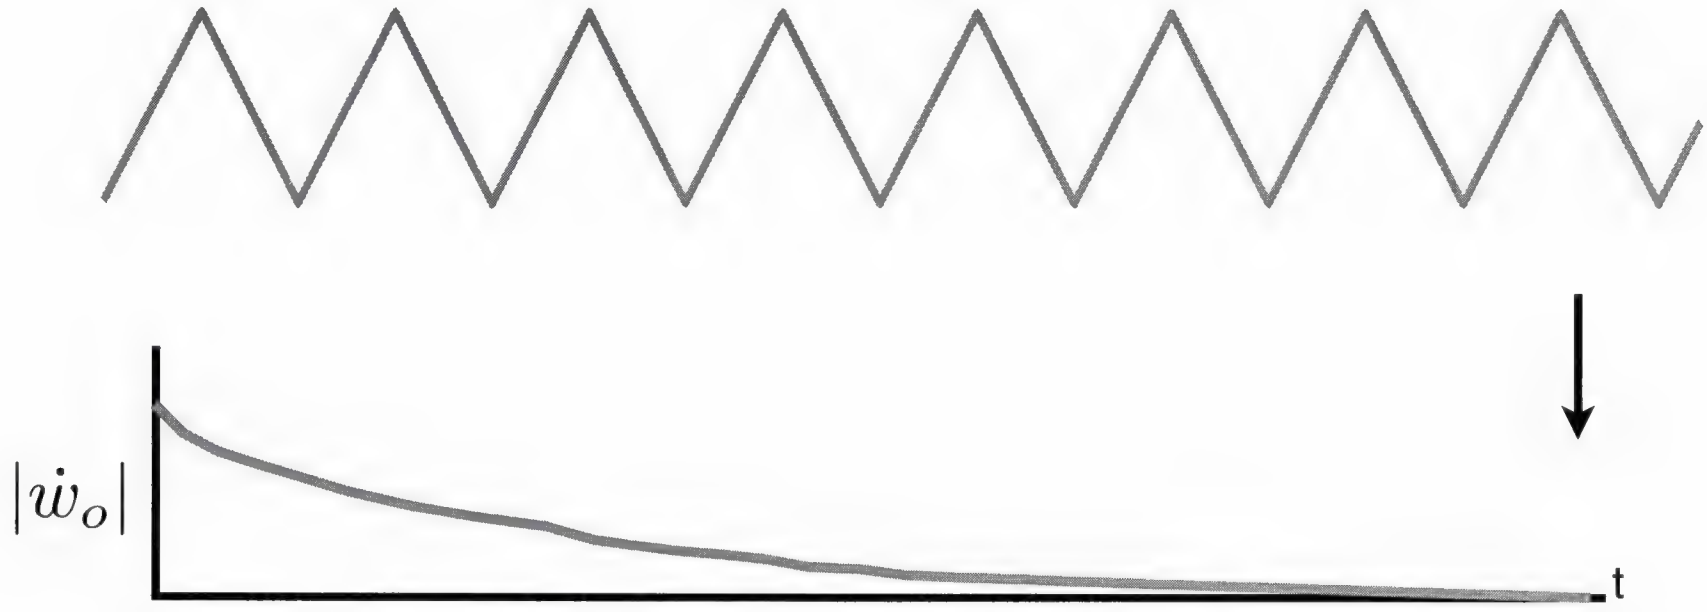
$|\dot{w}_o|$
t

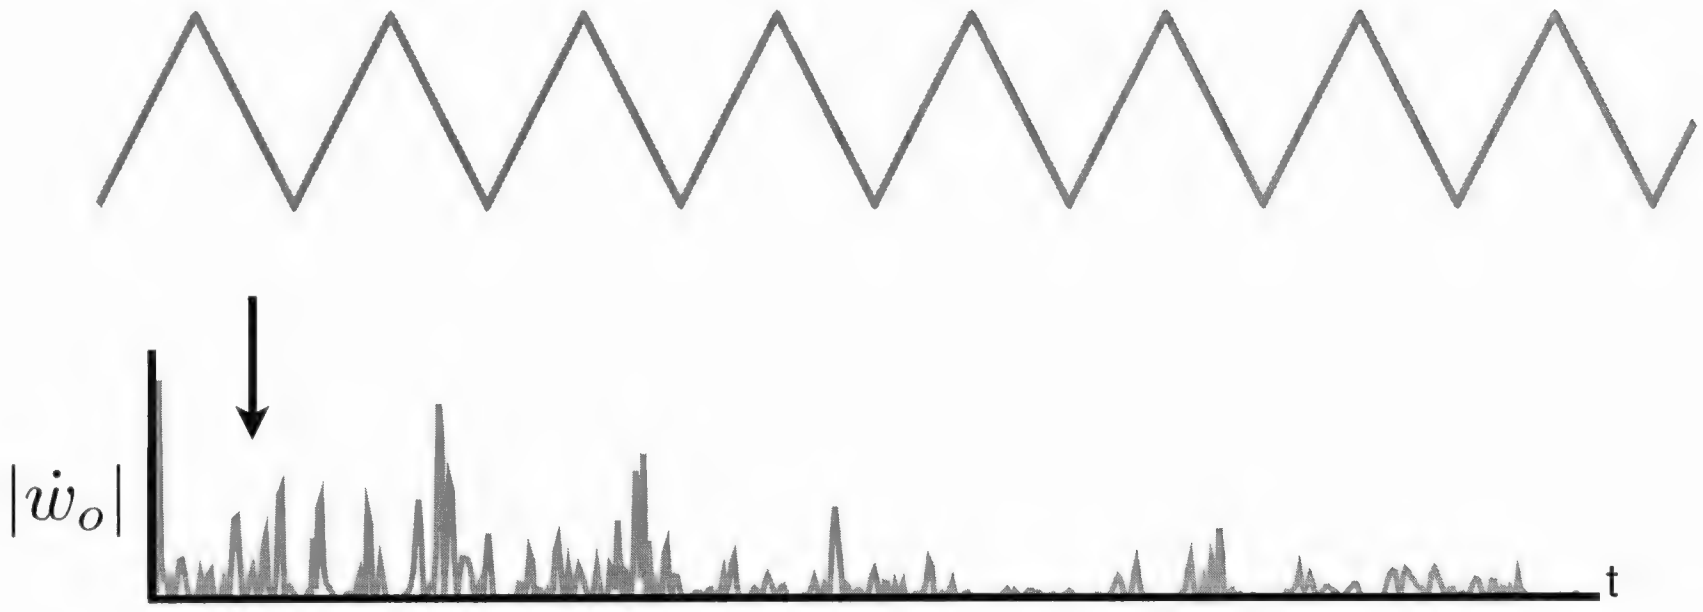
$|\dot{w}_o|$
t

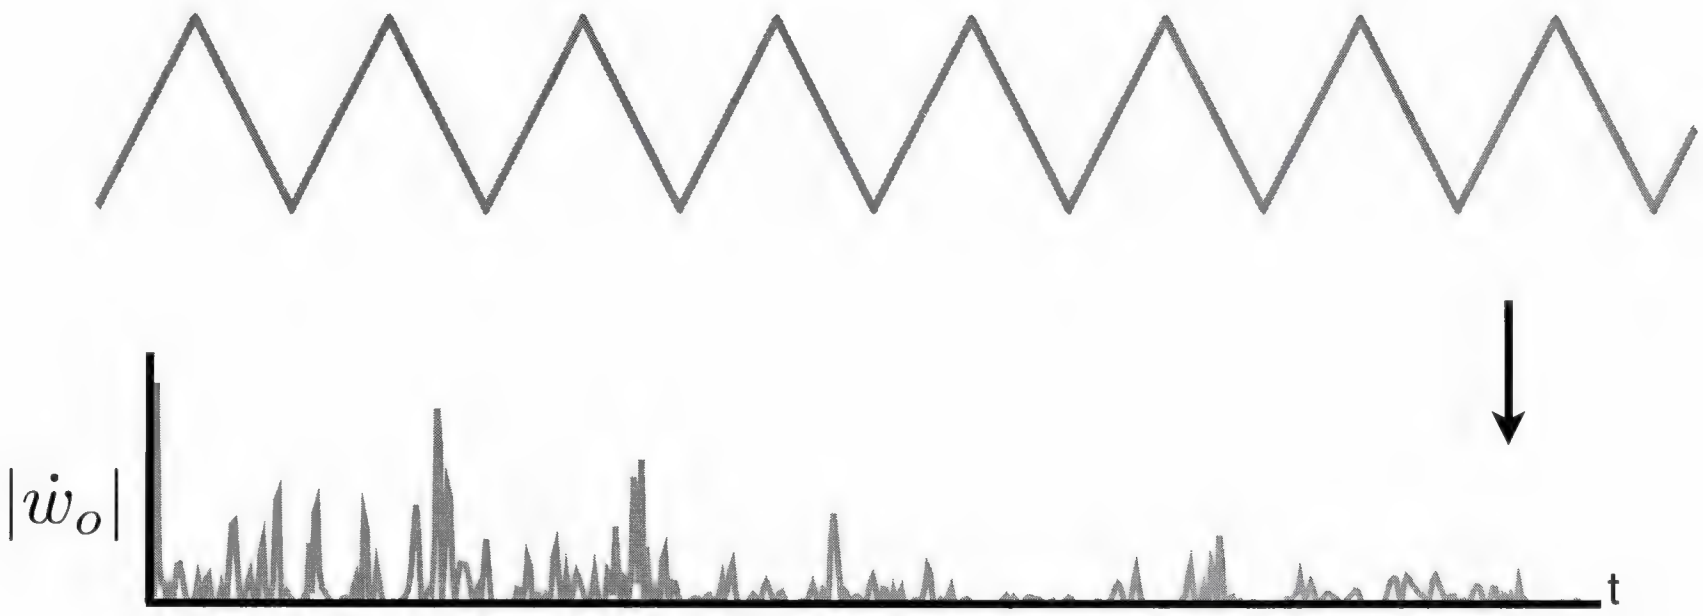
$|\dot{w}_o|$
t

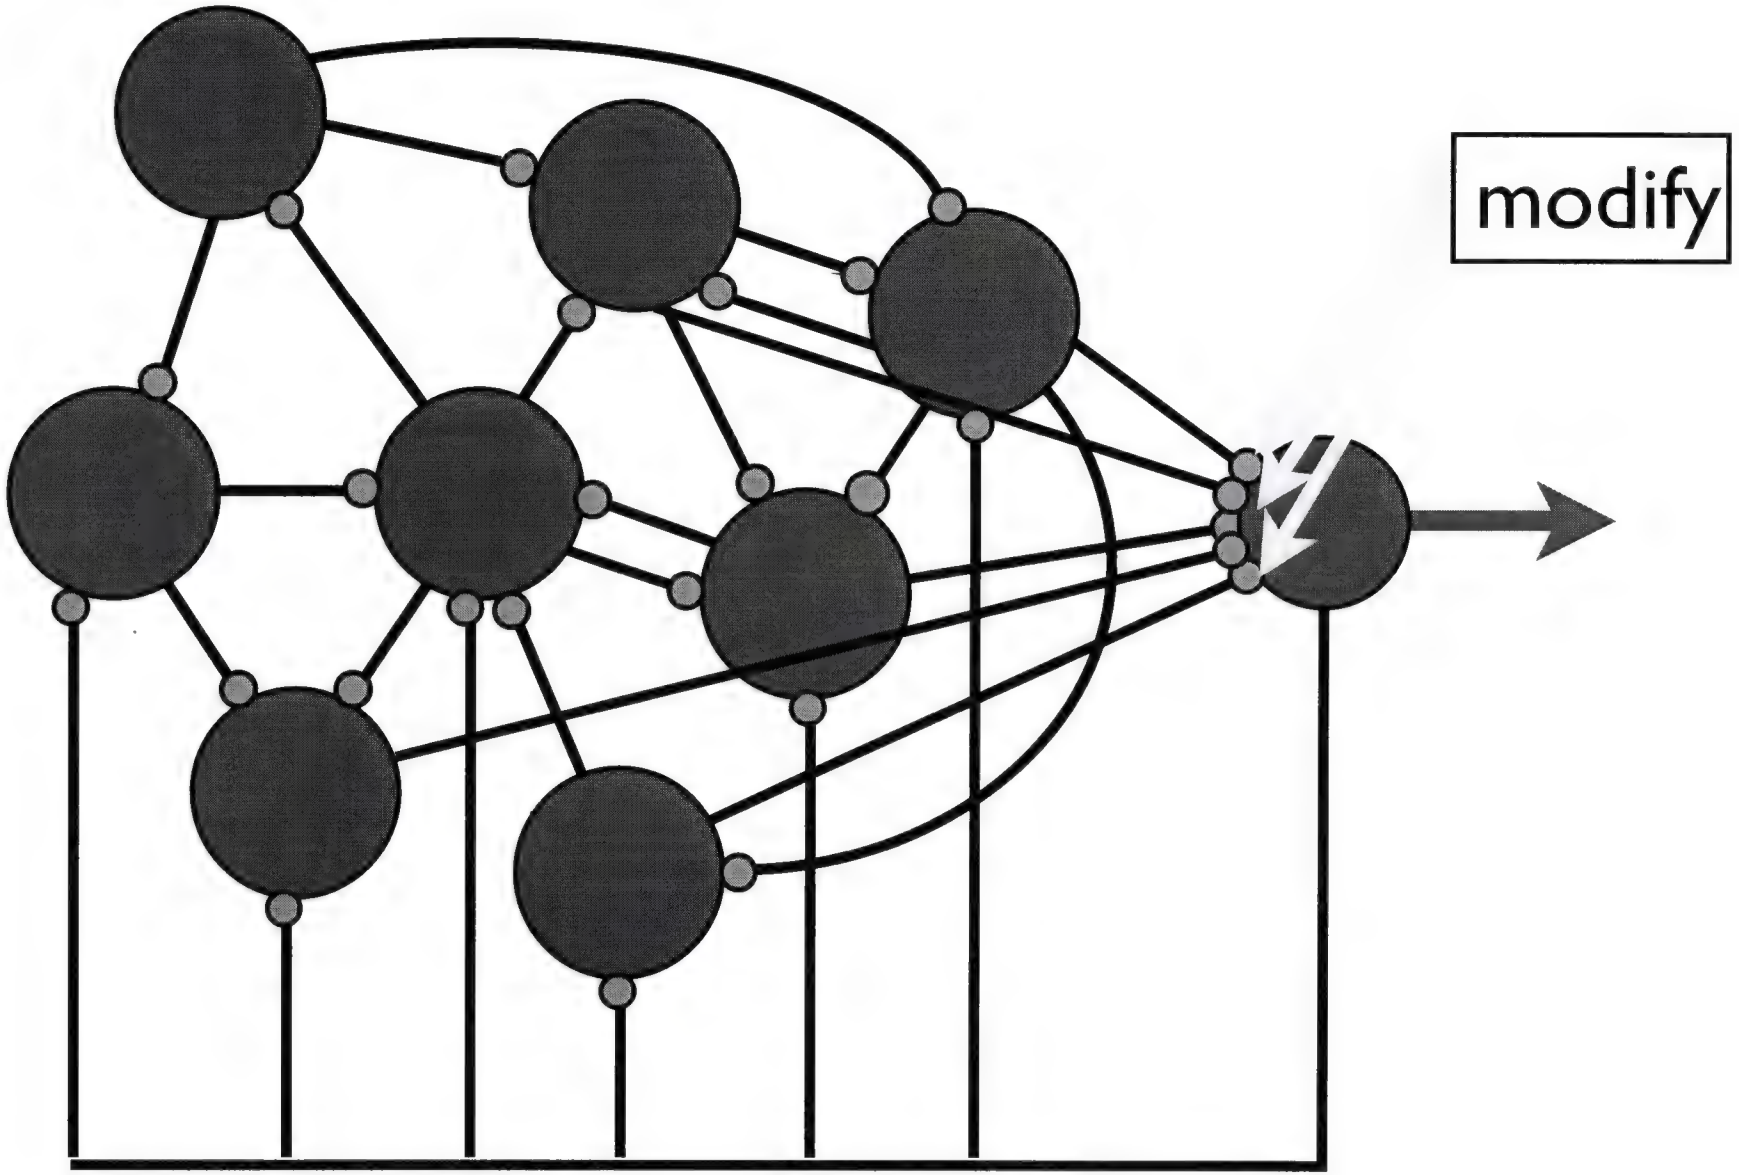
modify

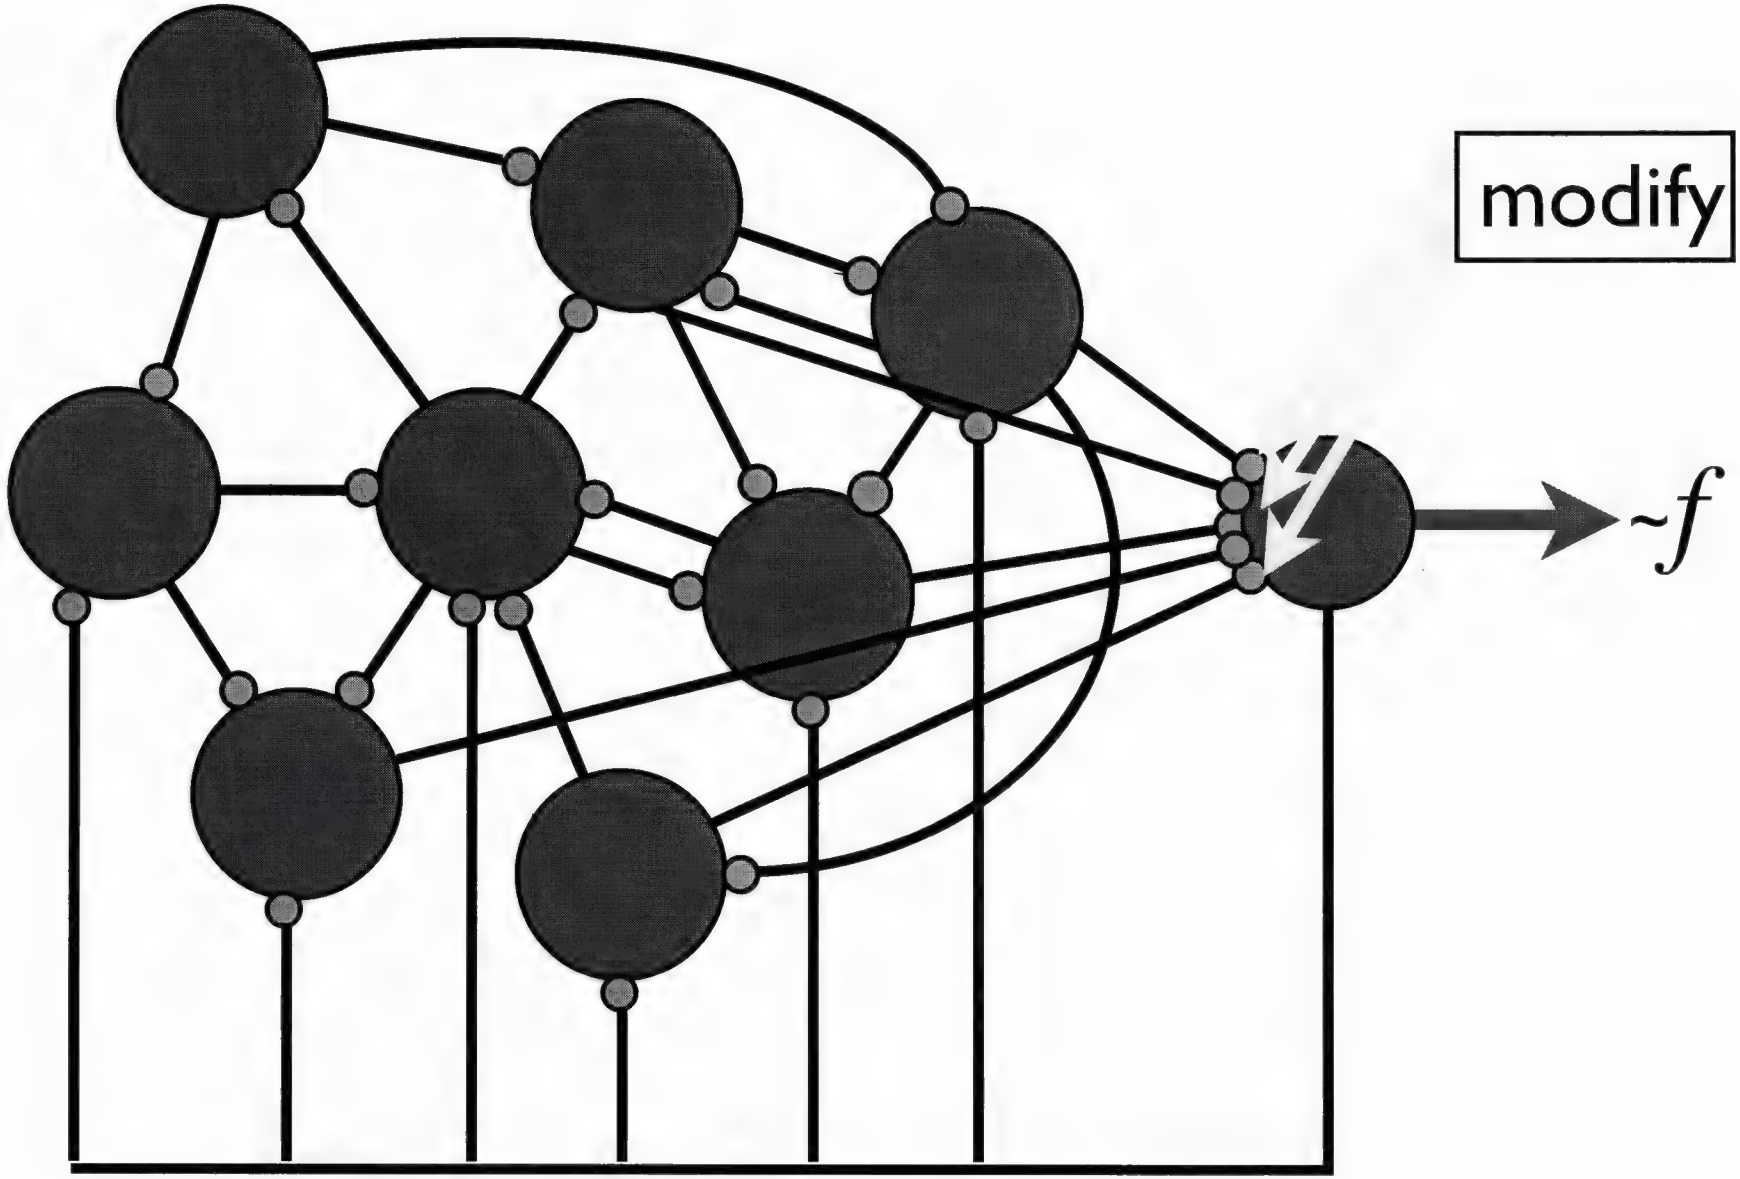
modify
~f

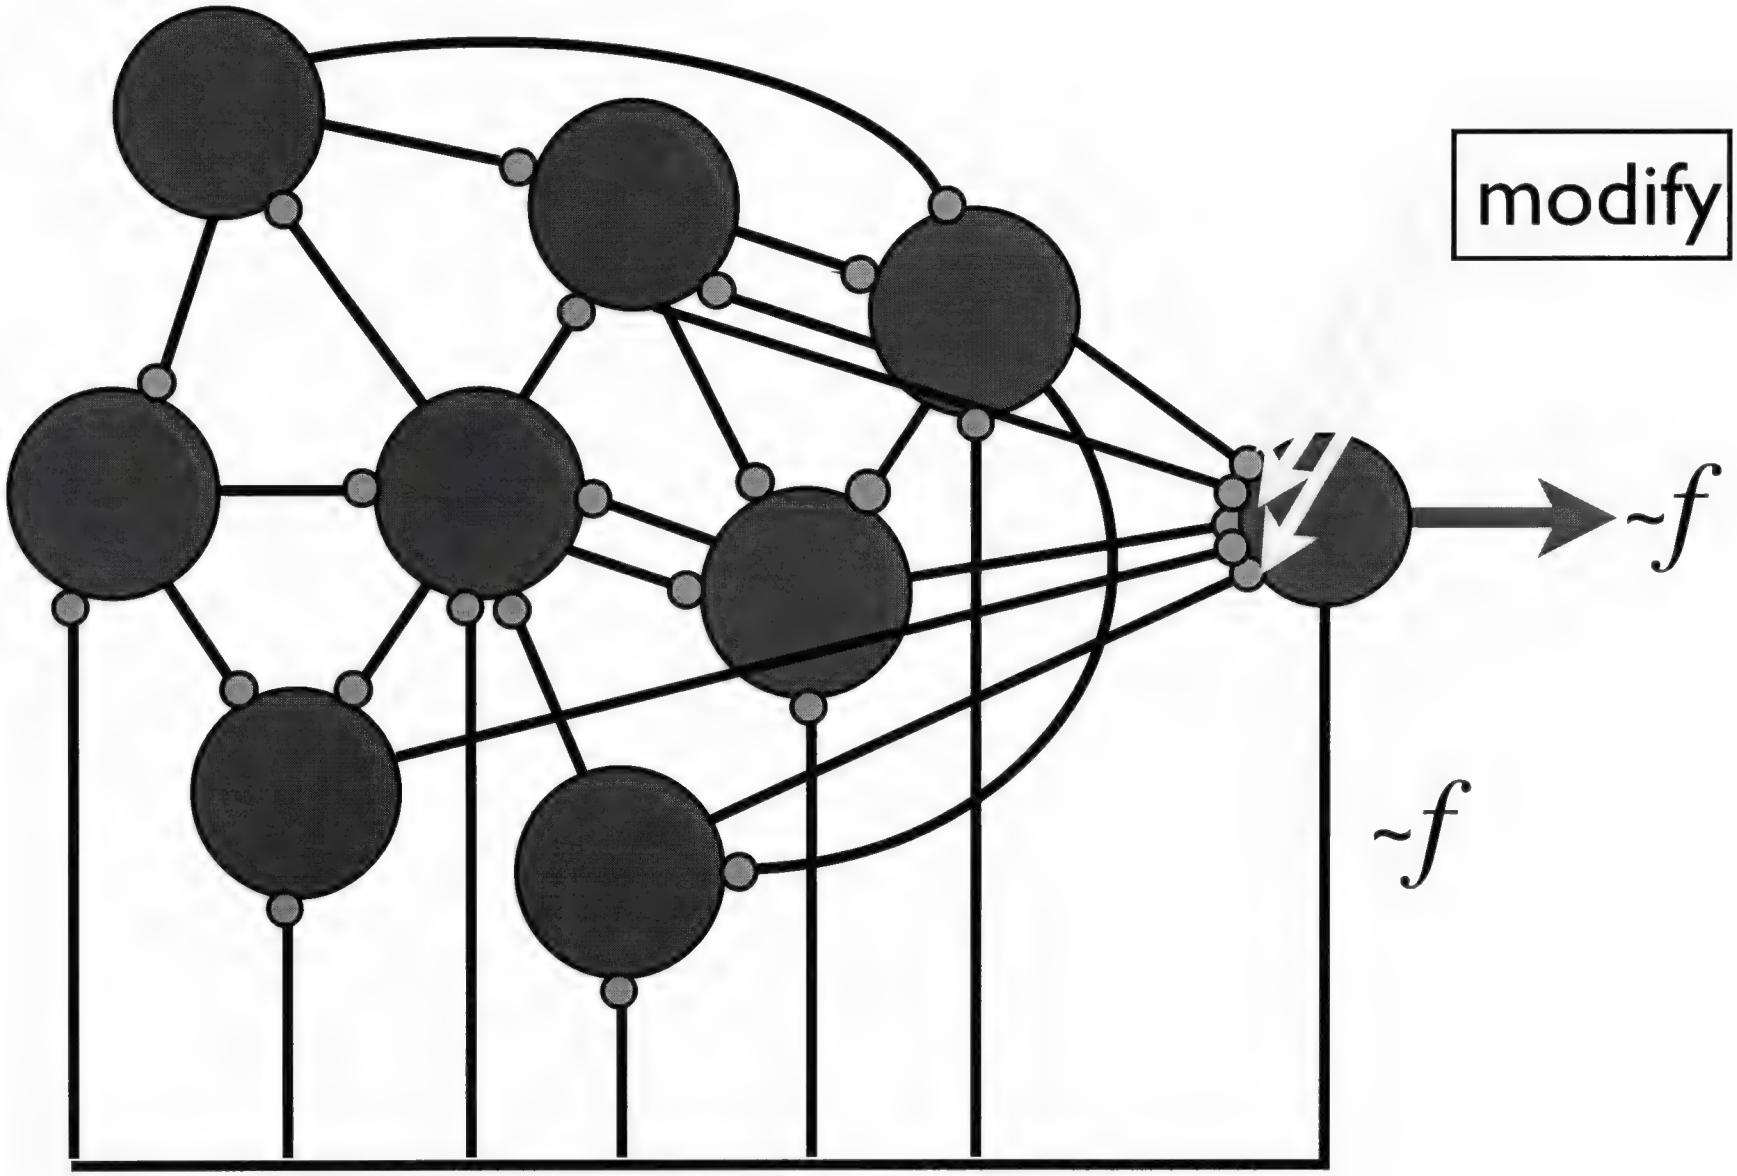
modify
~f
~f

First-Order Reduced & Controlled Error (FORCE) Learning

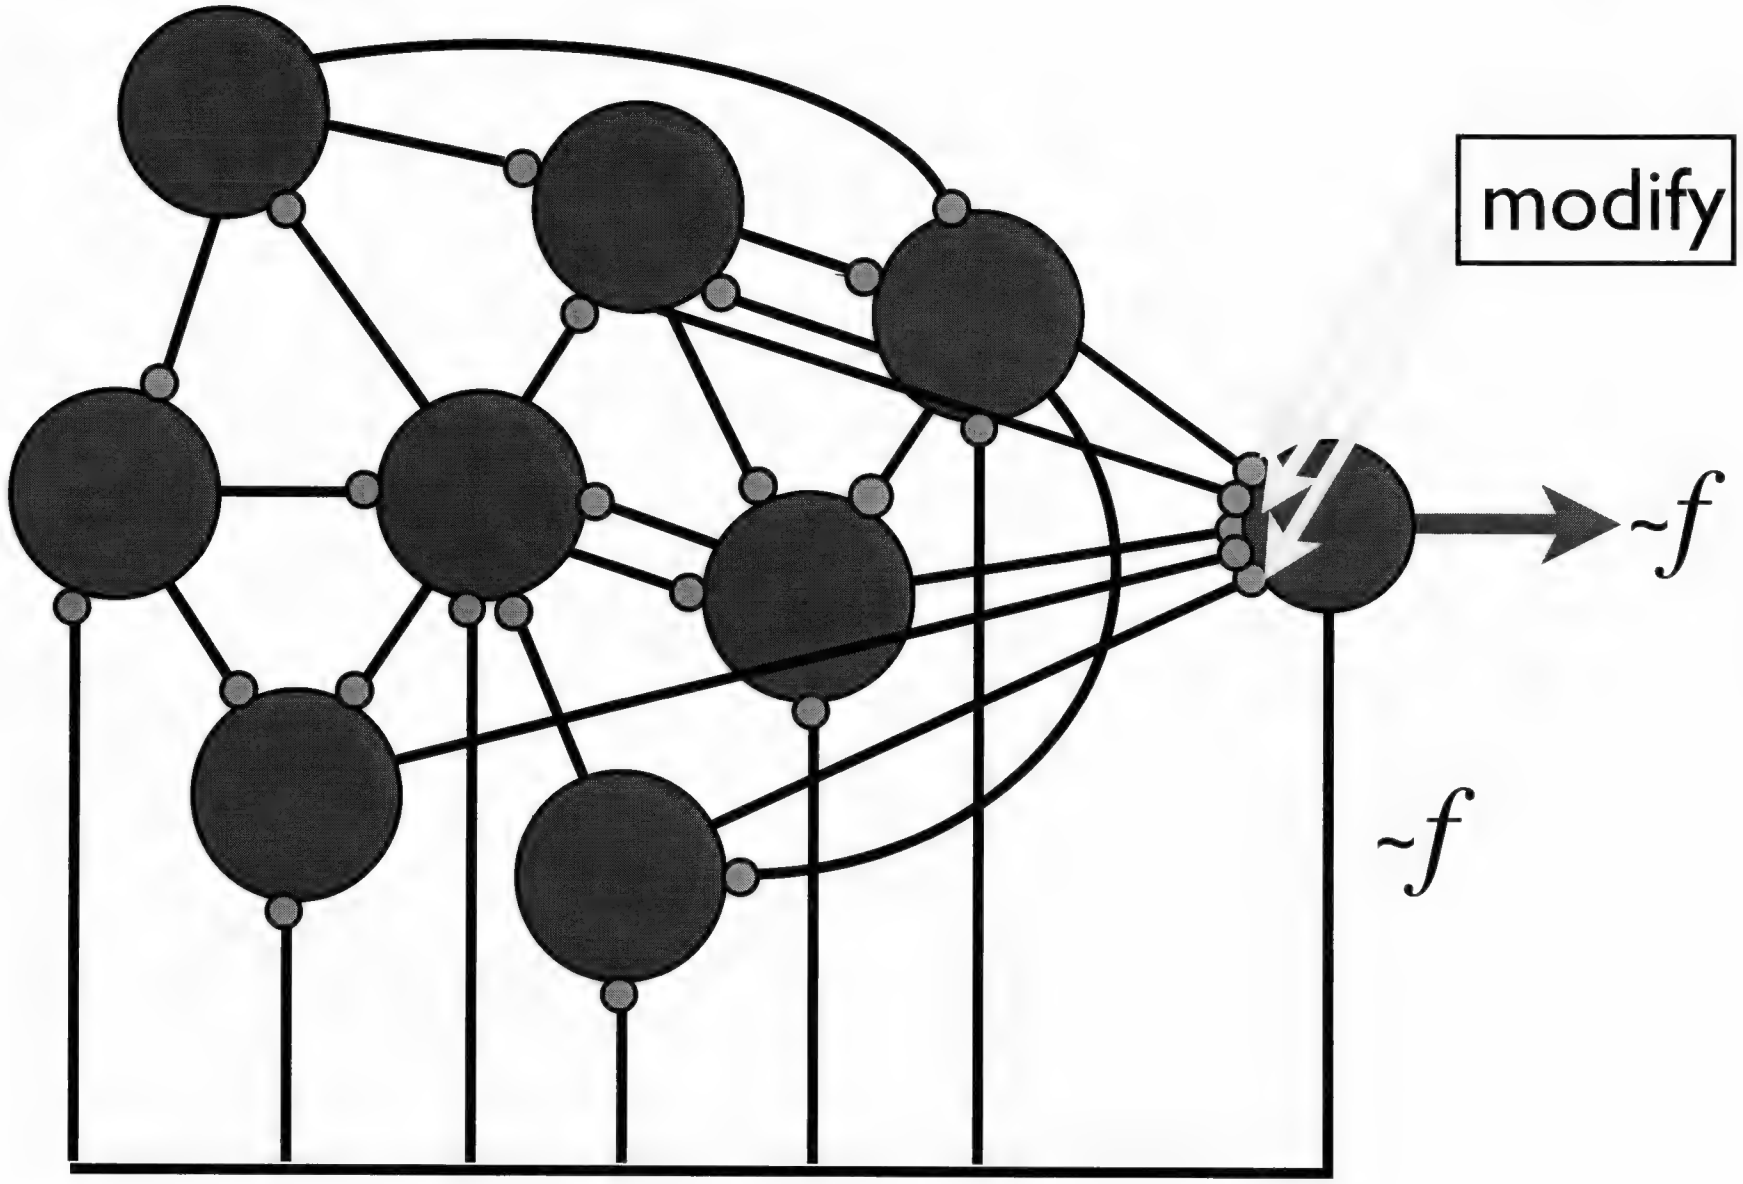

First-Order Reduced & Controlled Error (FORCE) Learning

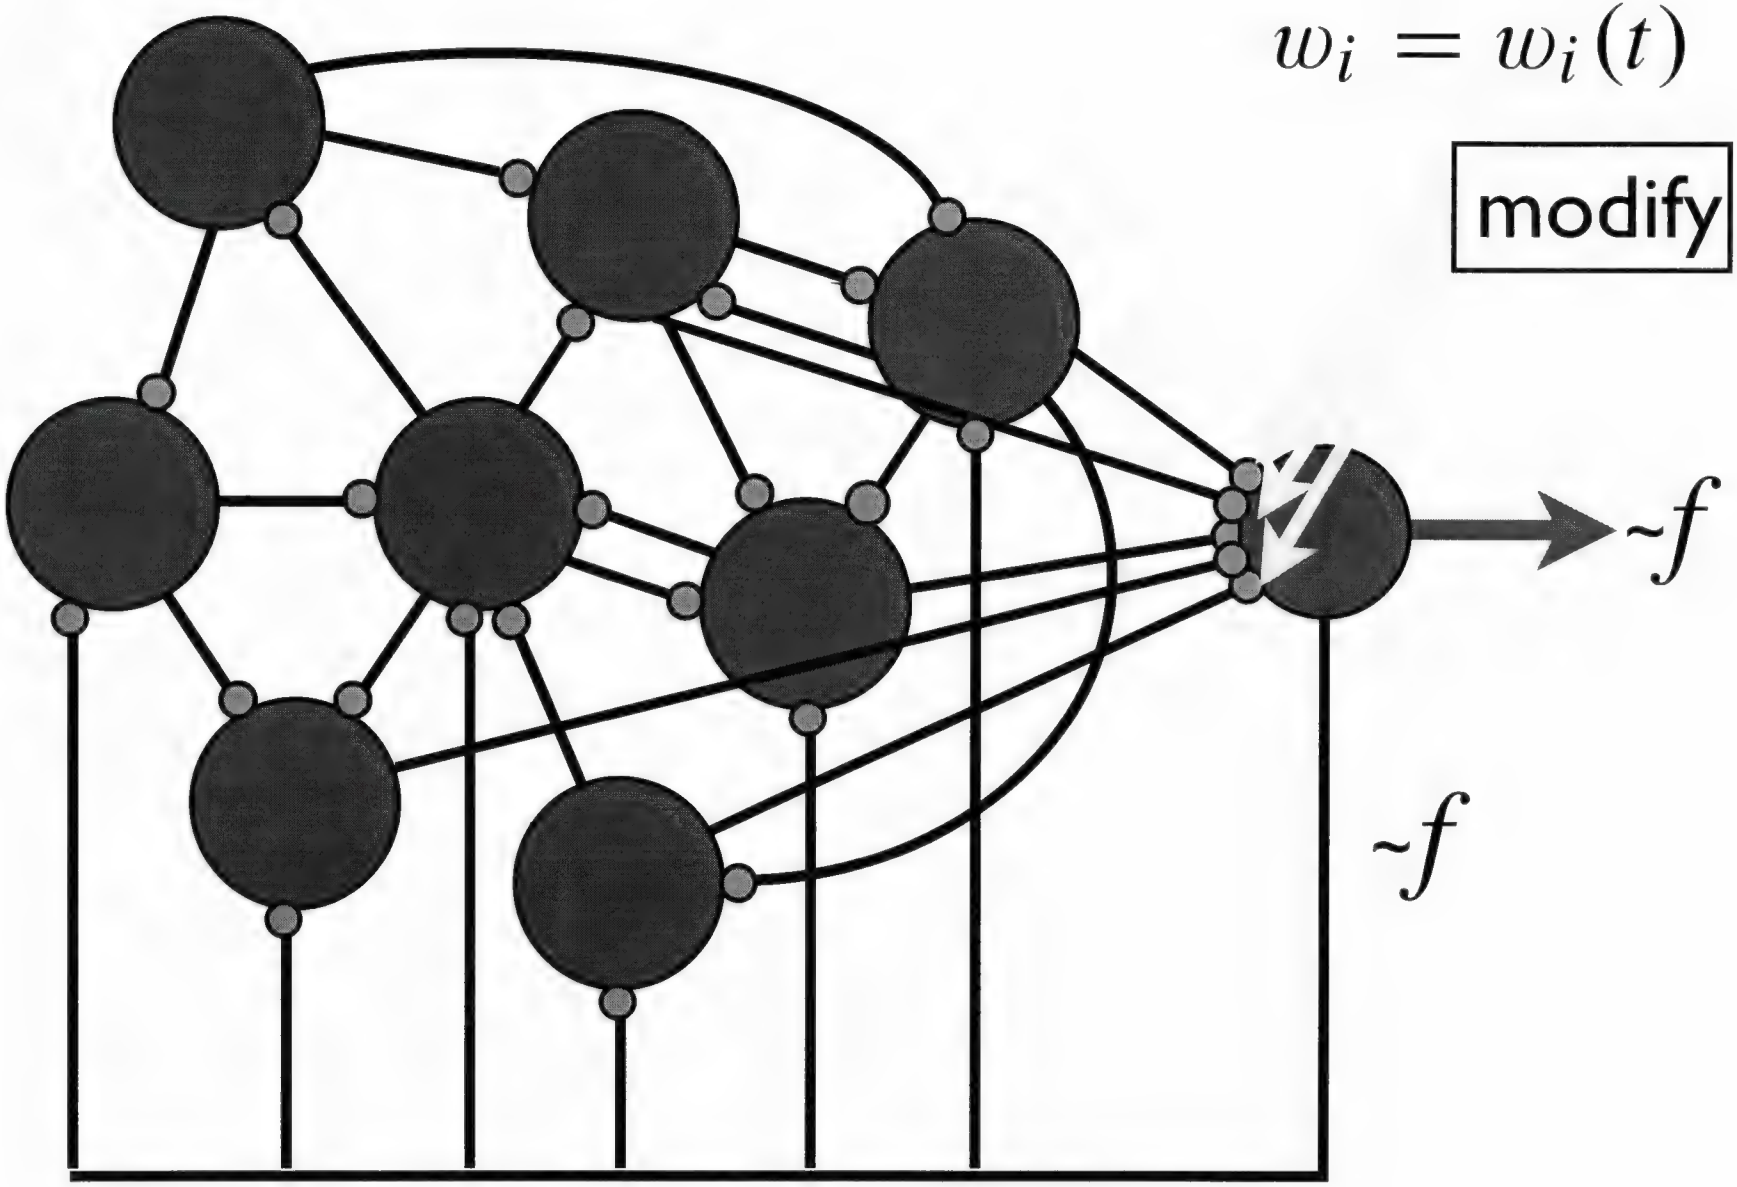

First-Order Reduced & Controlled Error (FORCE) Learning

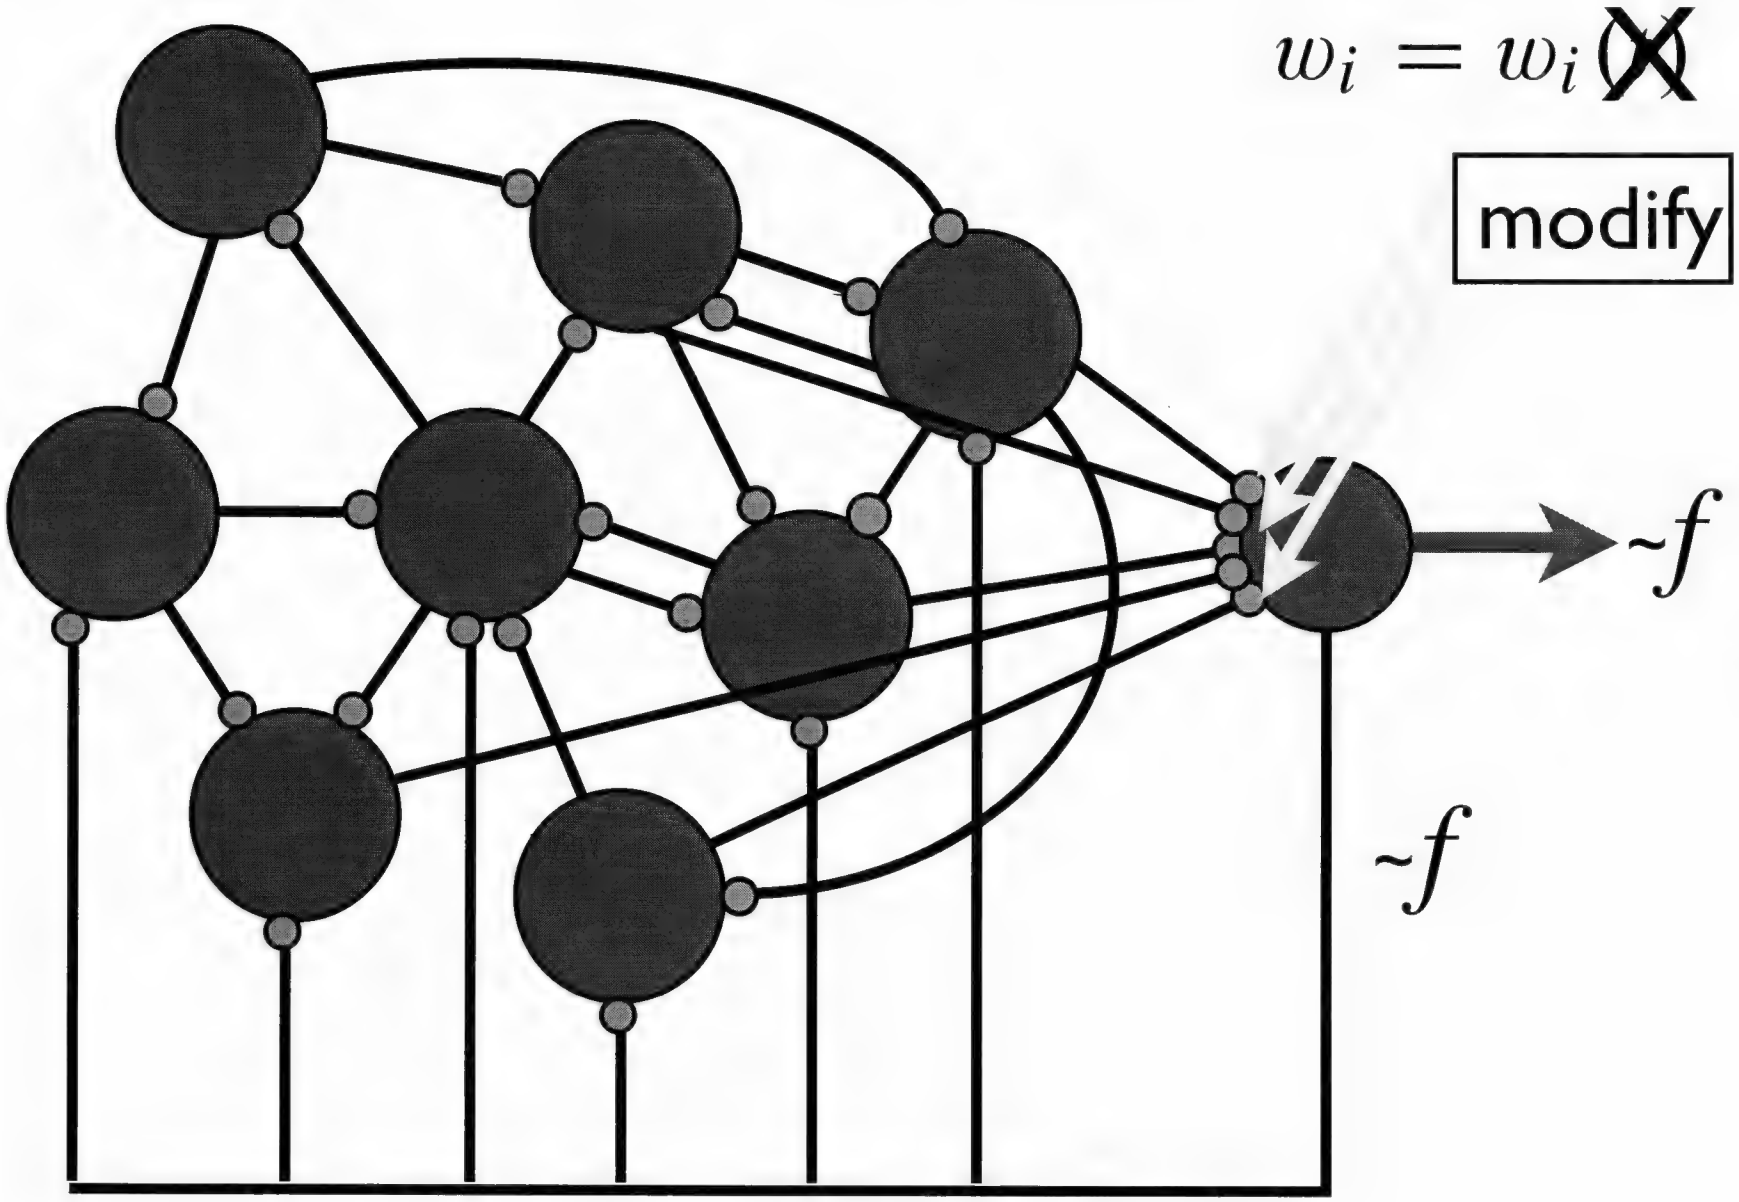

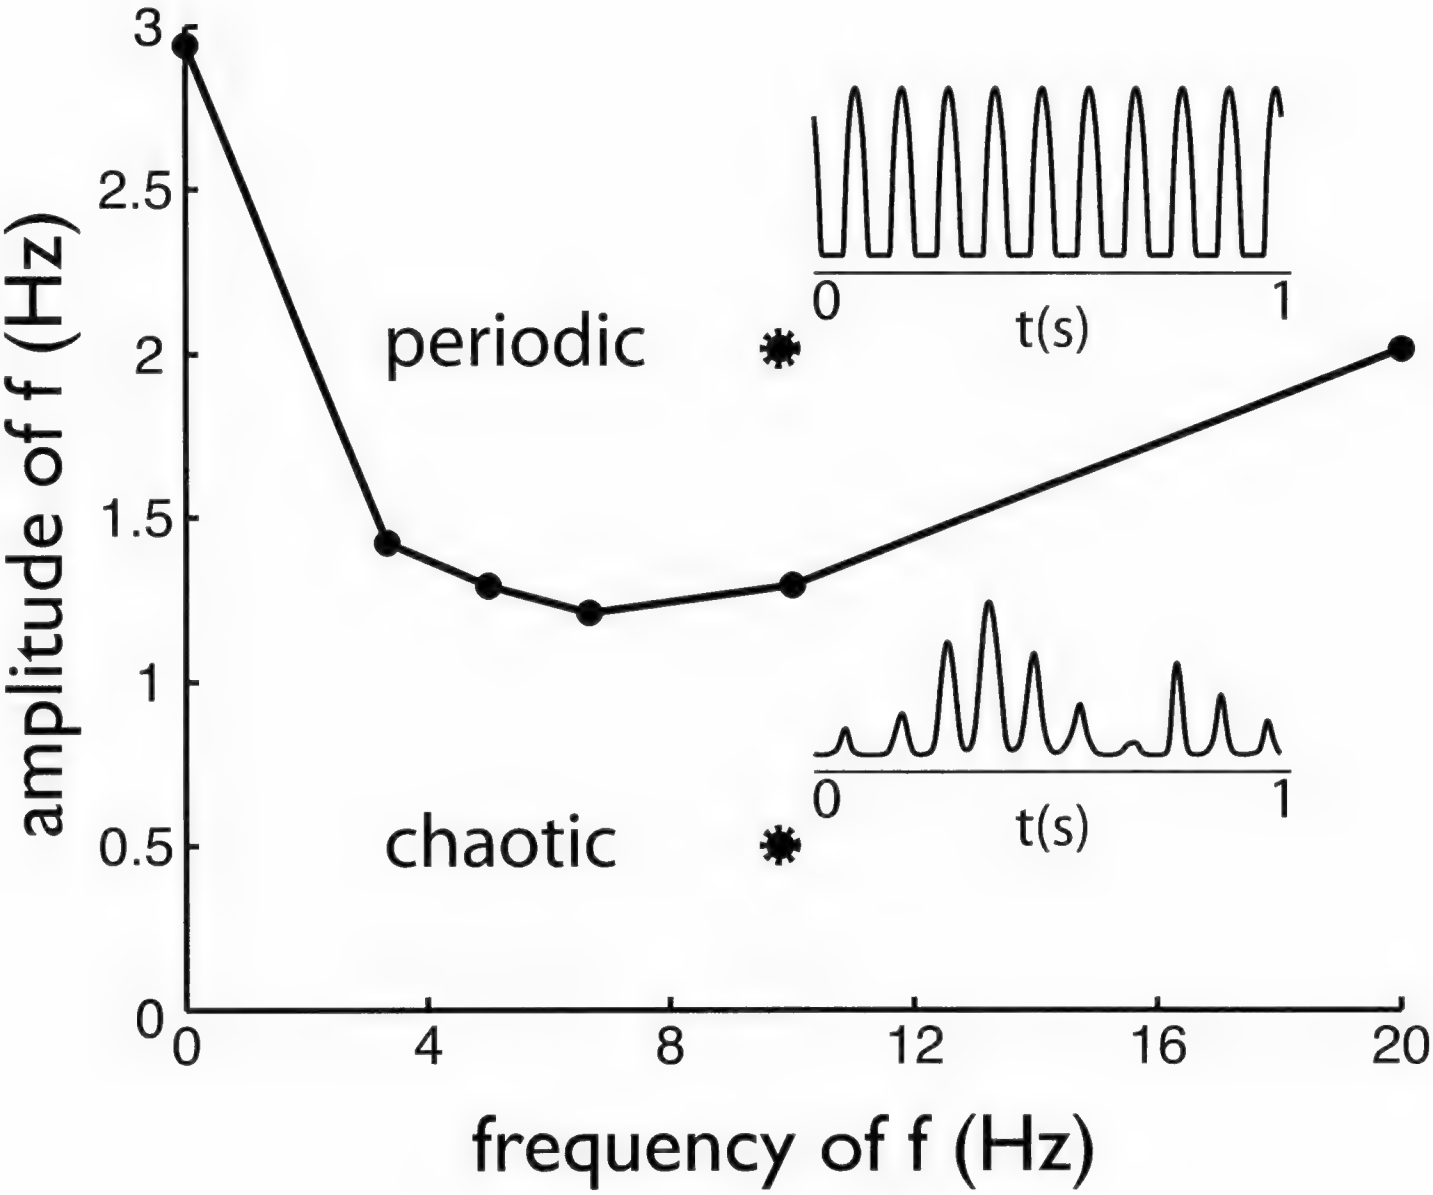

K. Rajan, H. Sompolinsky

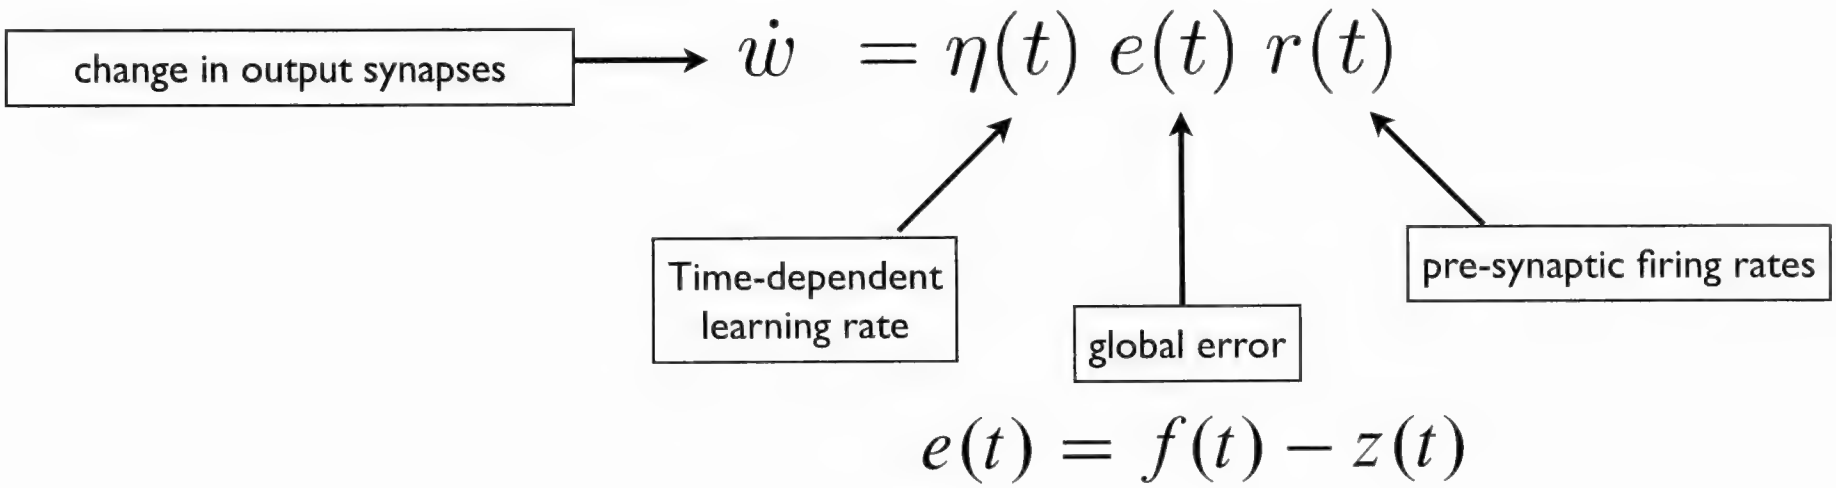
change in output synapses
$\dot{w} = \eta(t)\, e(t)\, r(t)$
Time-dependent learning rate
global error
pre-synaptic firing rates
$e(t) = f(t) - z(t)$

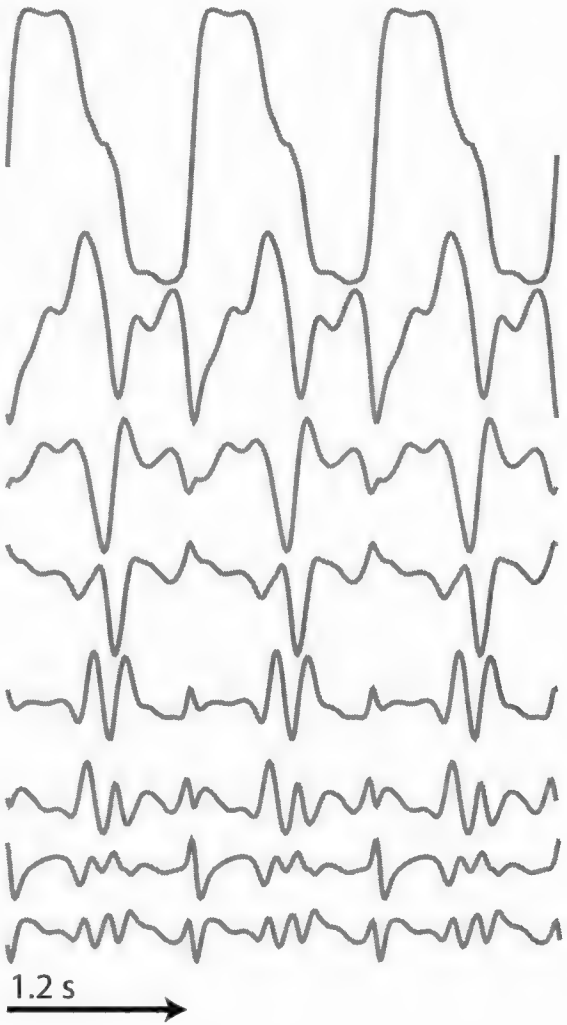
1.2 s

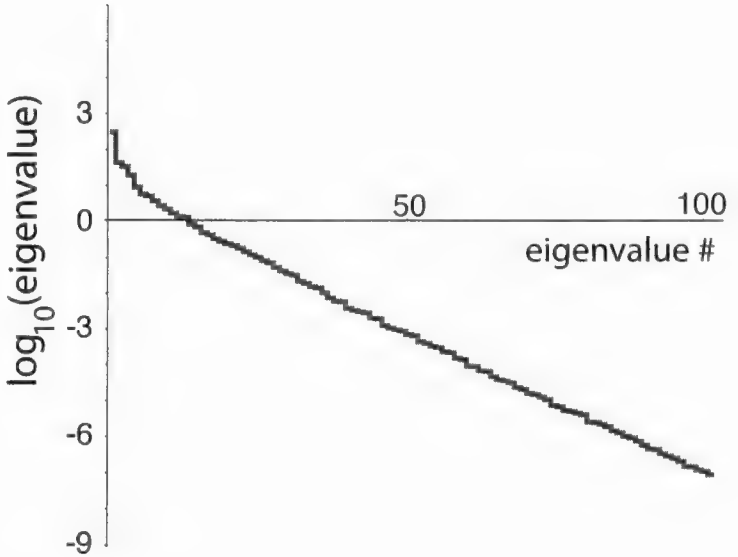
log10(eigenvalue)
3
0
-3
-6
-9
50
100
eigenvalue #

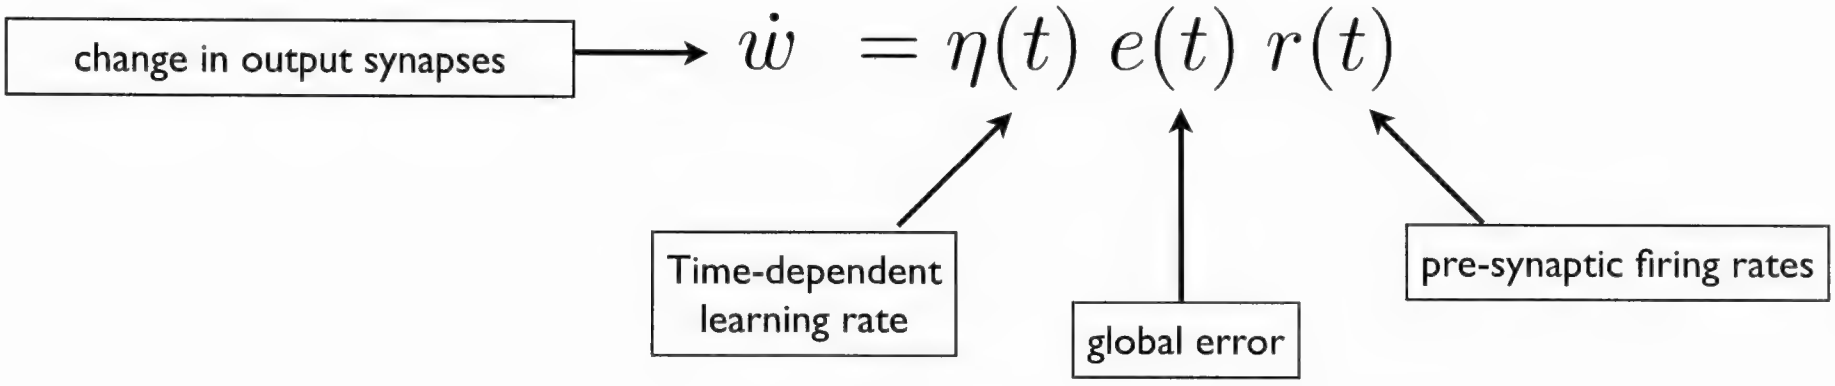

$$e(t) = f(t) - z(t)$$

$$\eta \propto \frac{1}{\lambda + \alpha}$$

change in output synapses

$$\dot{w} = \eta(t)\, e(t)\, r(t)$$

Time-dependent learning rate

global error

pre-synaptic firing rates

$$e(t) = f(t) - z(t)$$

$$\eta \propto \frac{1}{\lambda + \alpha}$$

$$\eta \to C^{-1} = \left(\alpha I + \langle r r^T \rangle\right)^{-1}$$

Recursive Least Squares Algorithm

$$\Delta \vec{w} = (f - z) P \cdot \vec{r}$$

$$P_{\text{initial}} = \alpha I \qquad \Delta P = - \frac{(P\vec{r})(\vec{r}^{\,T} P)}{1 + \vec{r}^{\,T} P \vec{r}}$$

Recursive Least Squares Algorithm

$$\Delta \vec{w} = (f - z) P \cdot \vec{r}$$

$$P_{\text{initial}} = \alpha I \qquad \Delta P = - \frac{(P\vec{r})(\vec{r}^{\,T} P)}{1 + \vec{r}^{\,T} P \vec{r}}$$

$$\vec{w}(0) = 0$$

after one time step

$$z = \vec{w} \cdot \vec{r} = f \left(\frac{\alpha |r|^2}{1 + \alpha |r|^2} \right) \rightarrow f$$

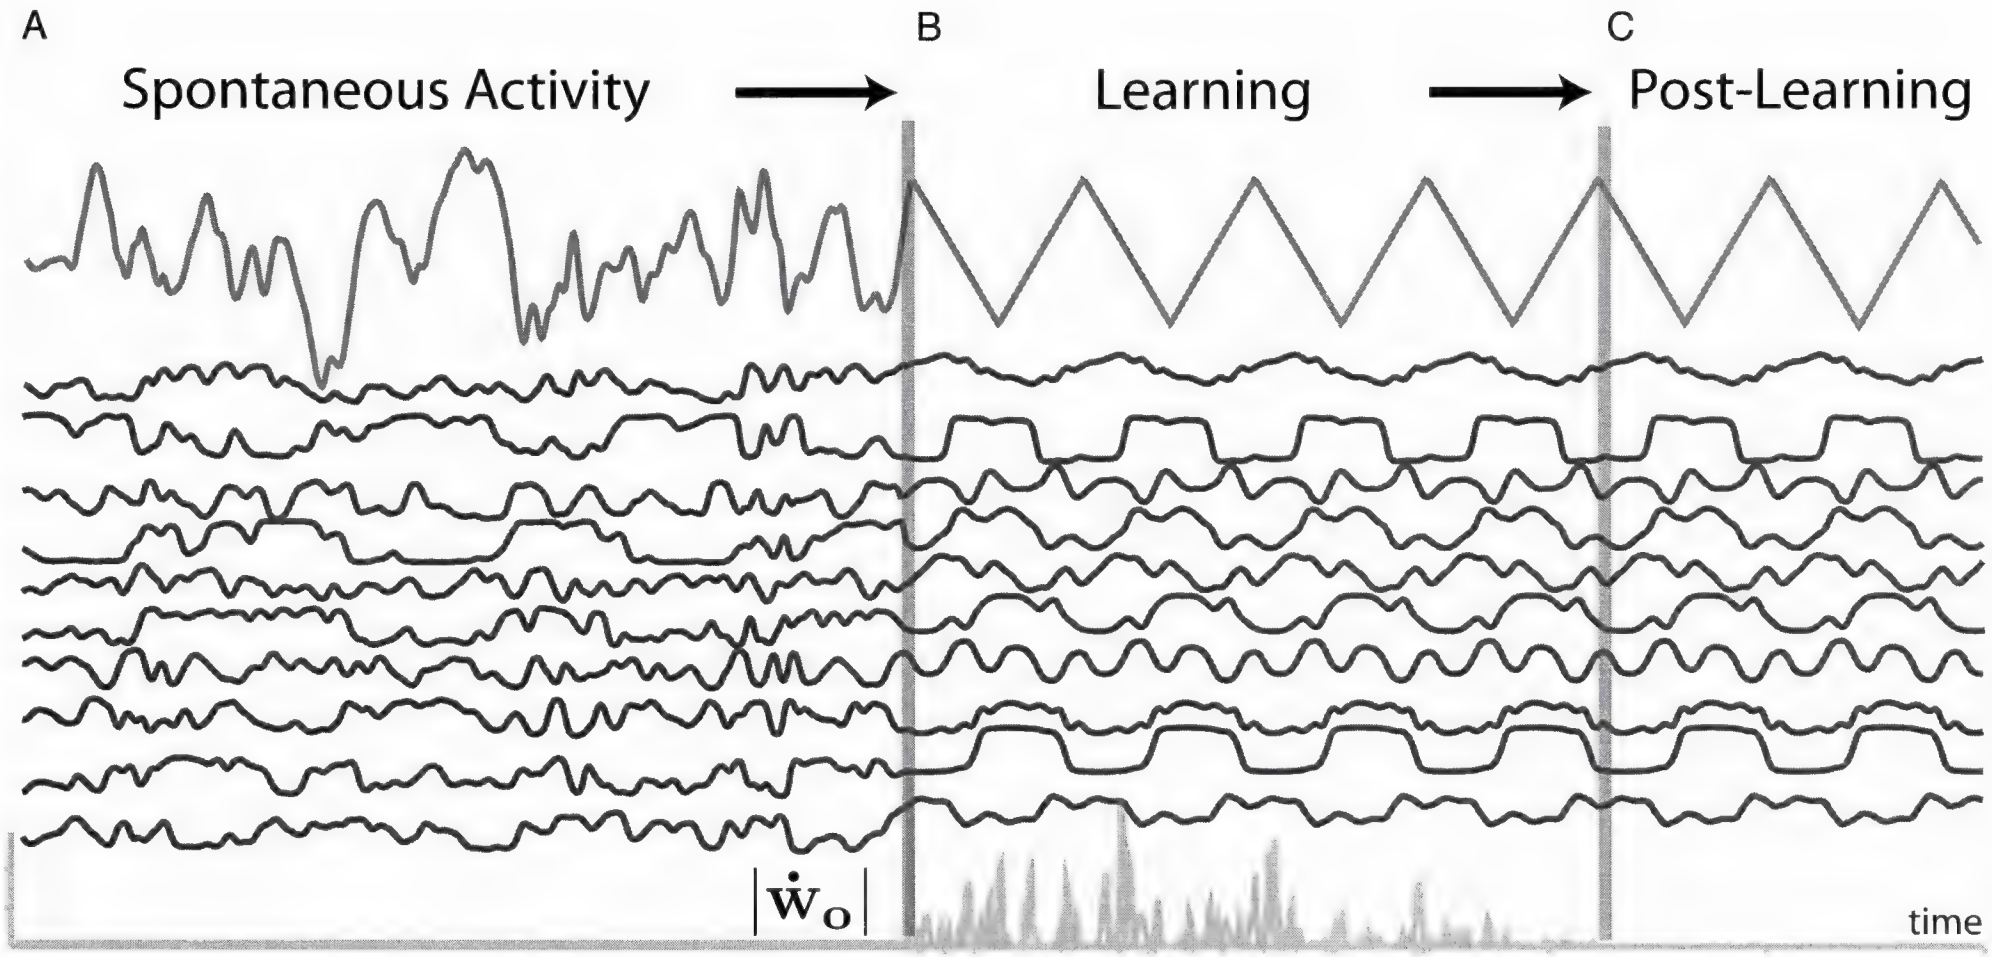
A
B
C
Spontaneous Activity
Learning
Post-Learning
$|\dot{\mathbf{w}}_o|$
time

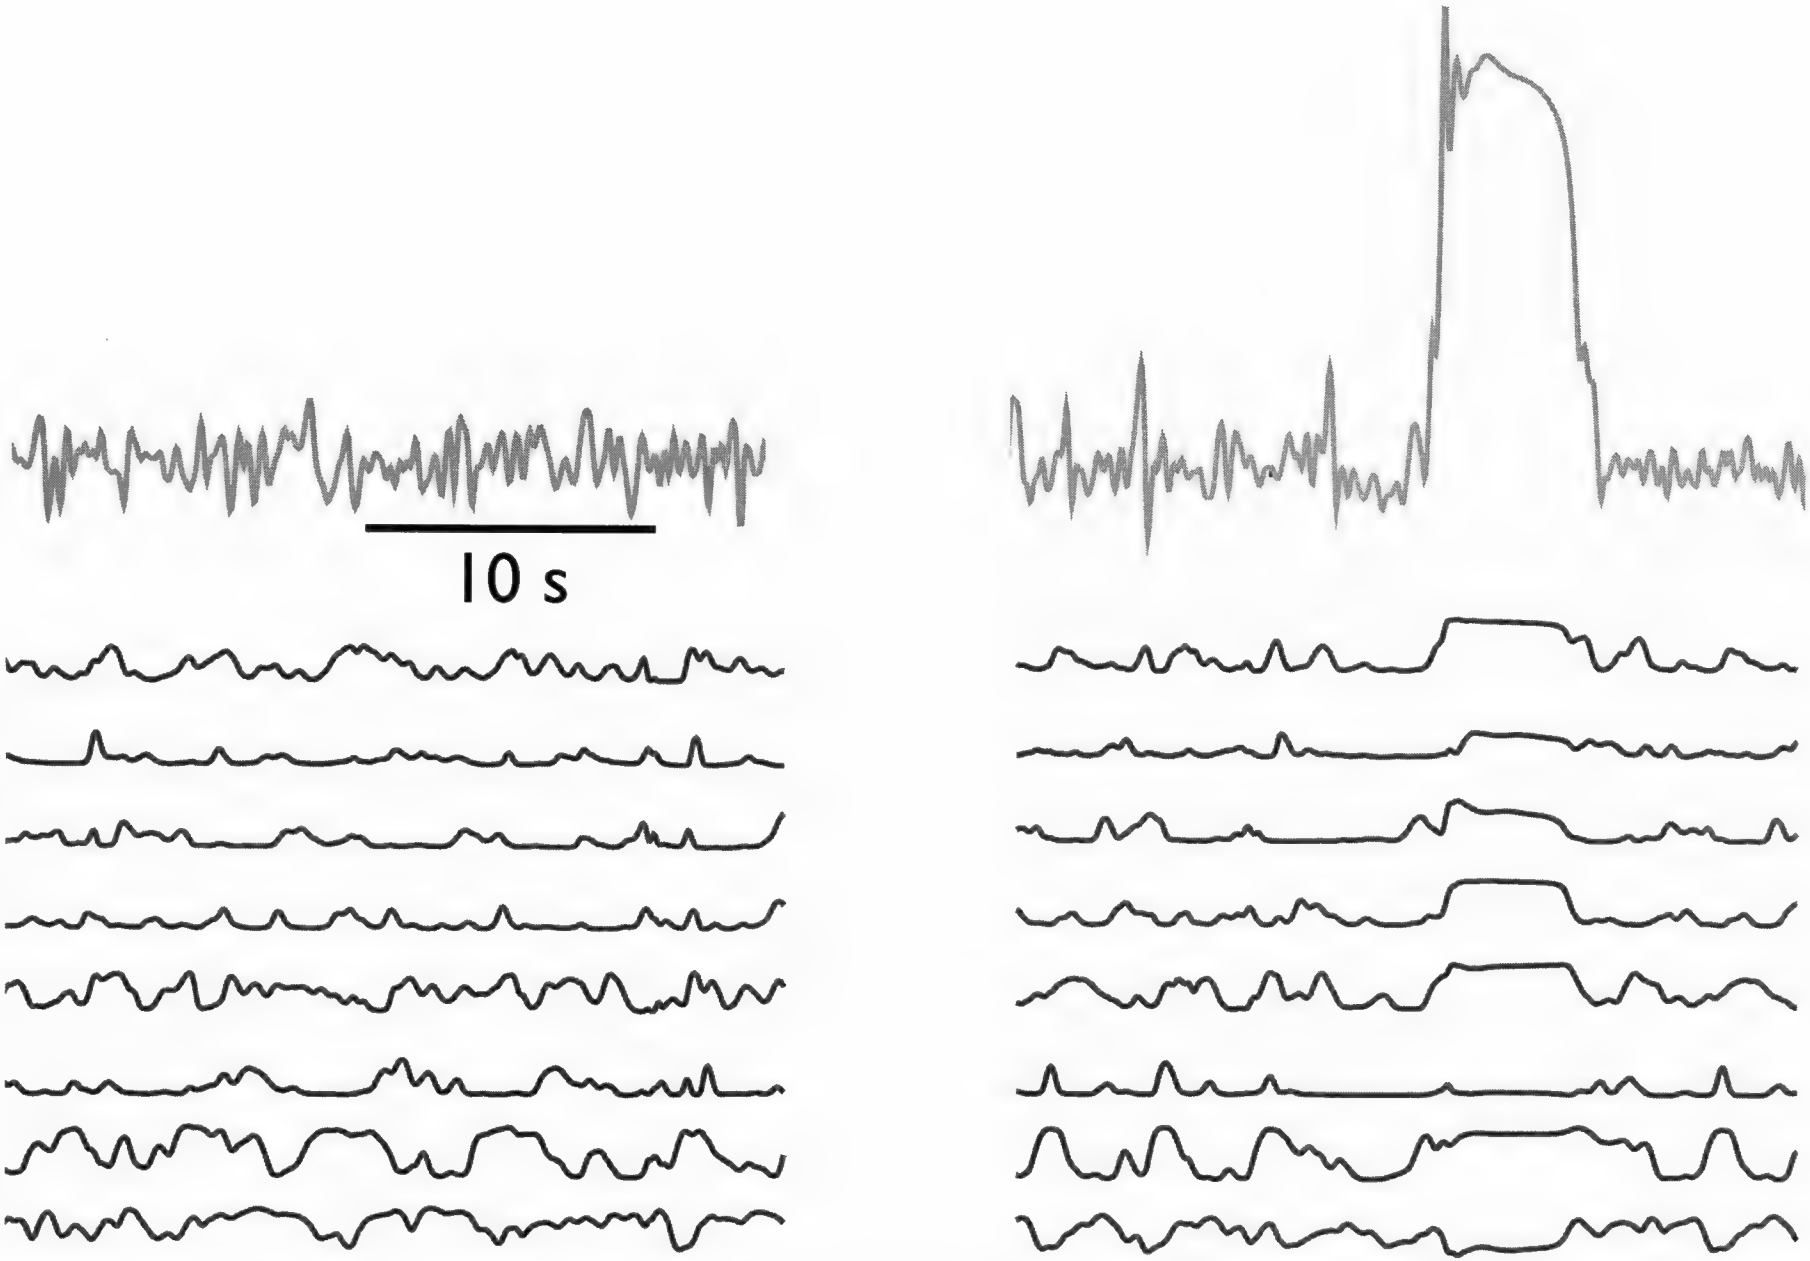
10 s

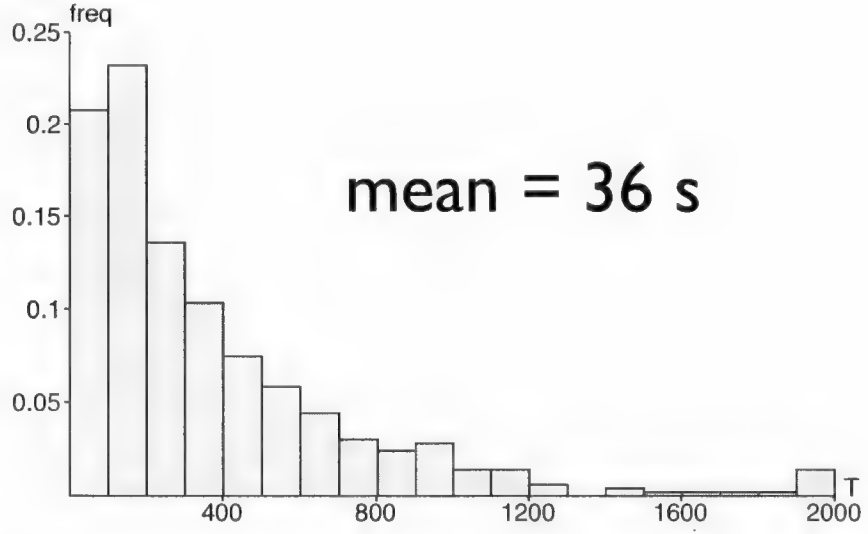

mean = 36 s

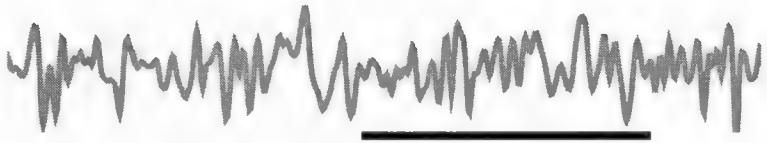

10 s

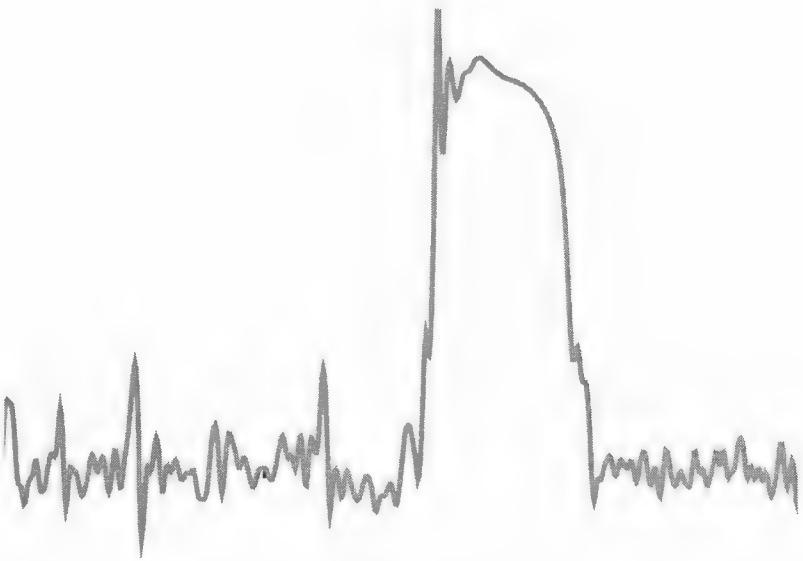

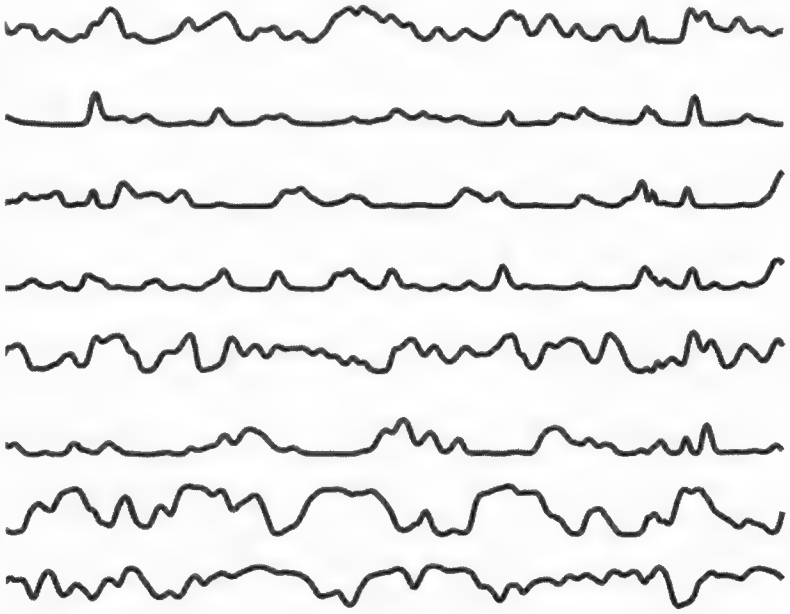

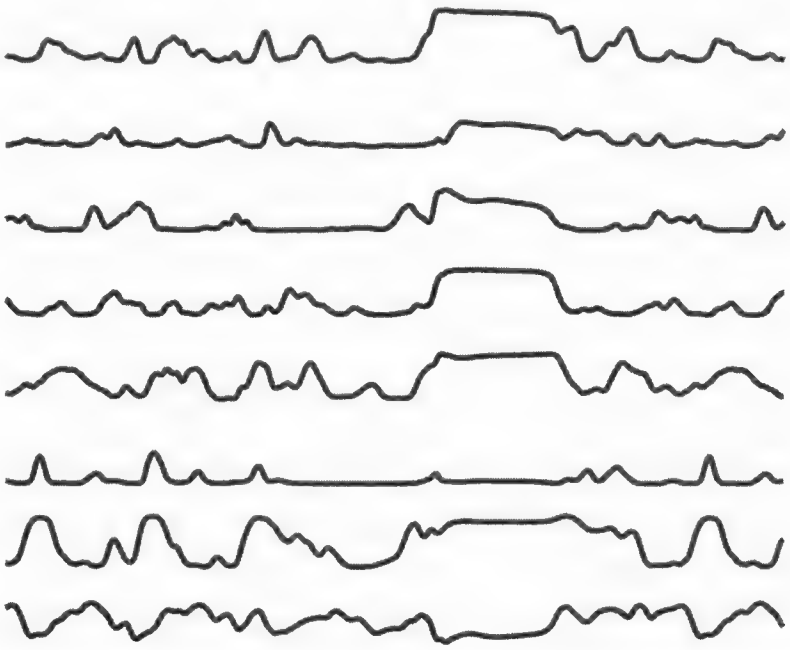

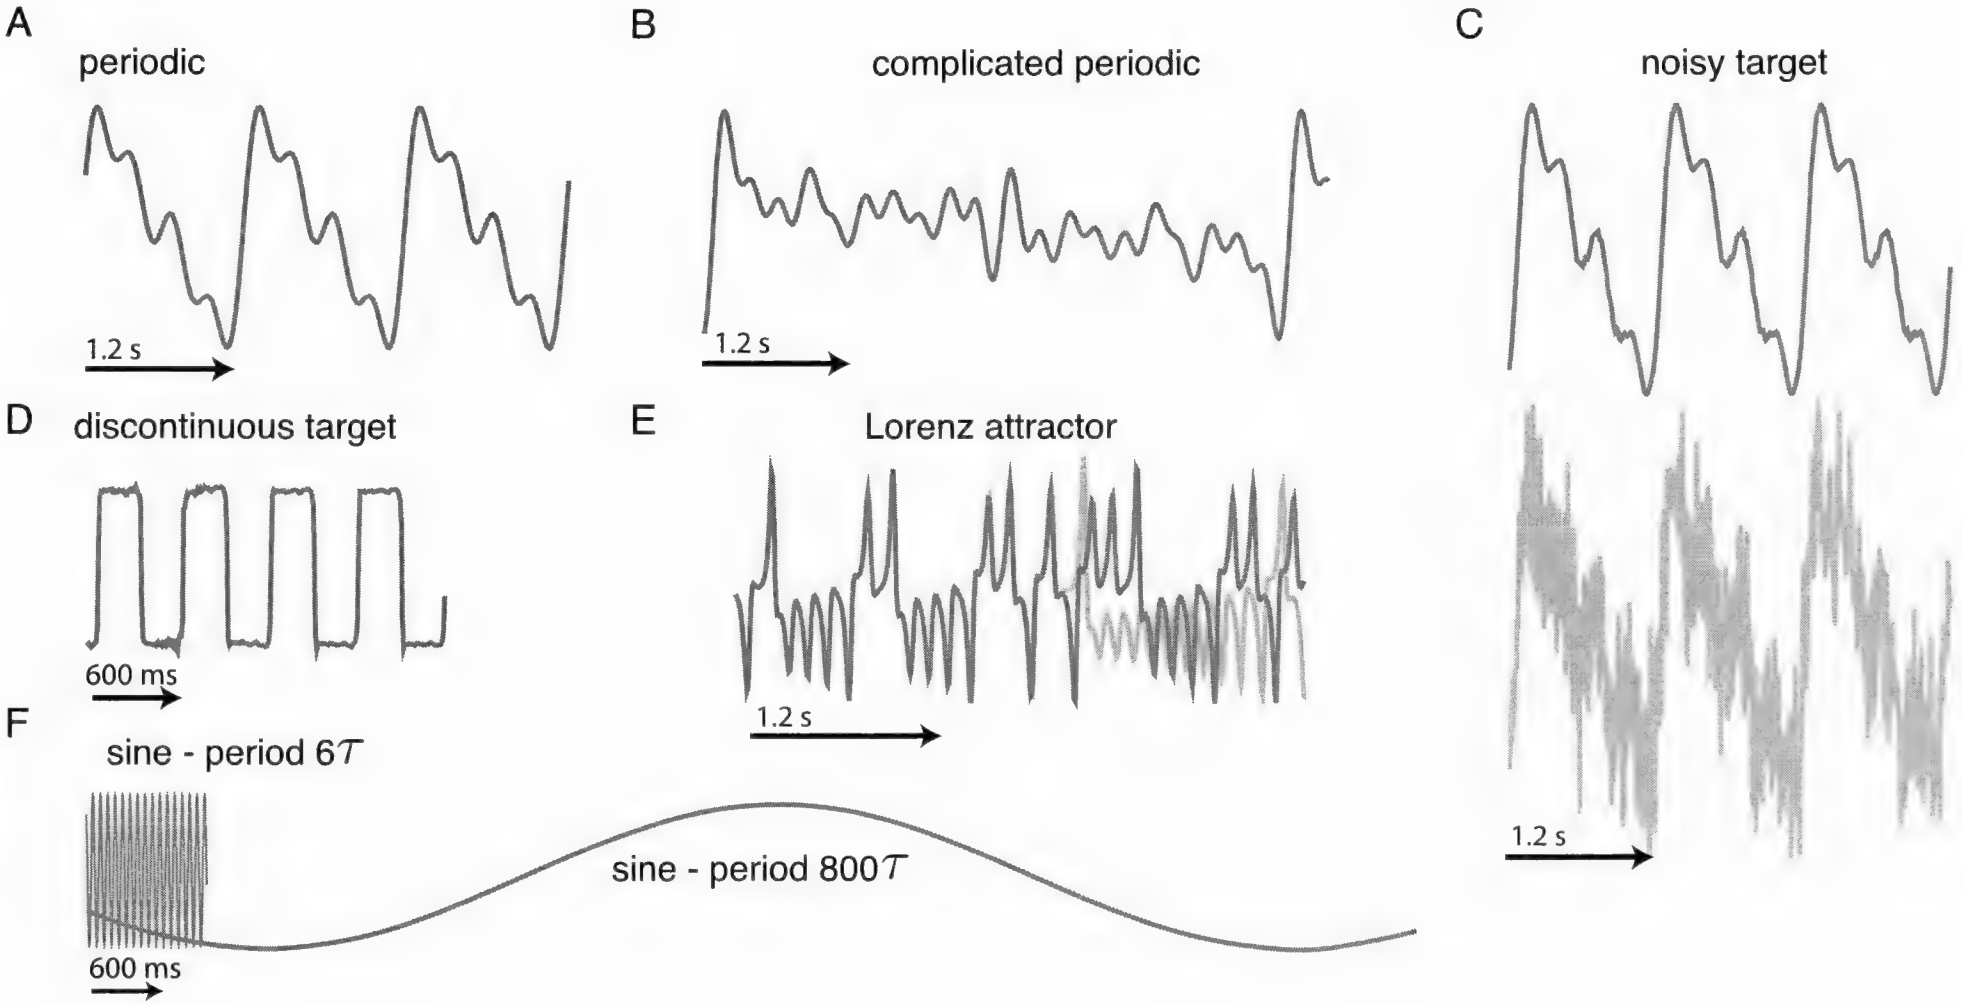
A
periodic
1.2 s
B
complicated periodic
1.2 s
C
noisy target
1.2 s
D
discontinuous target
600 ms
E
Lorenz attractor
1.2 s
F
sine - period 6$\mathcal{T}$
sine - period 800$\mathcal{T}$
600 ms

Aperiodic Example

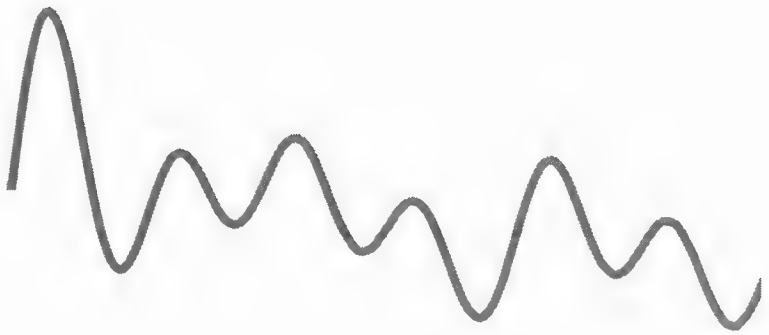

Aperiodic Example

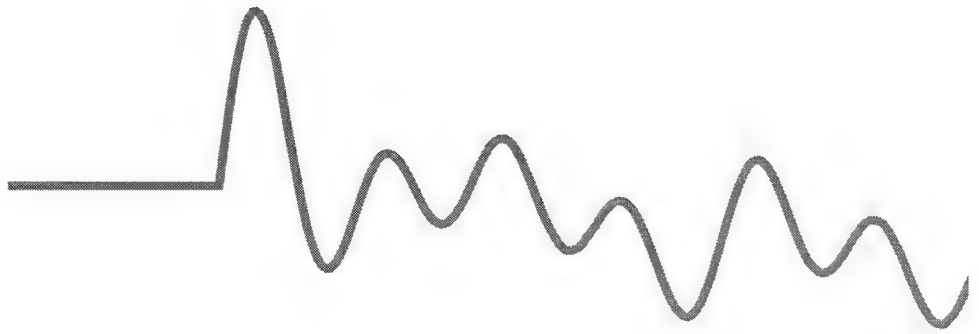

Aperiodic Example

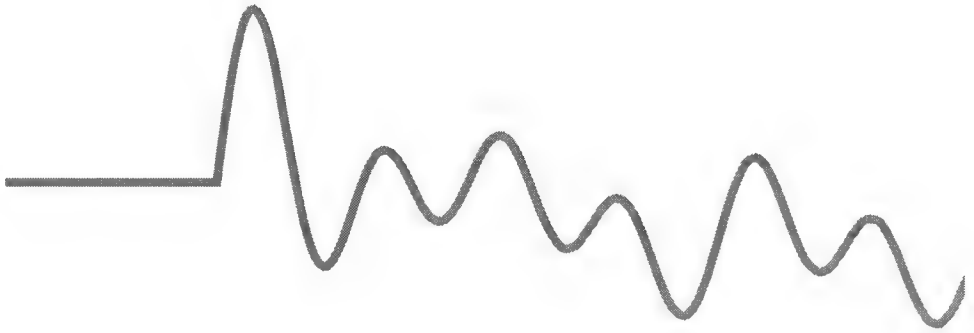

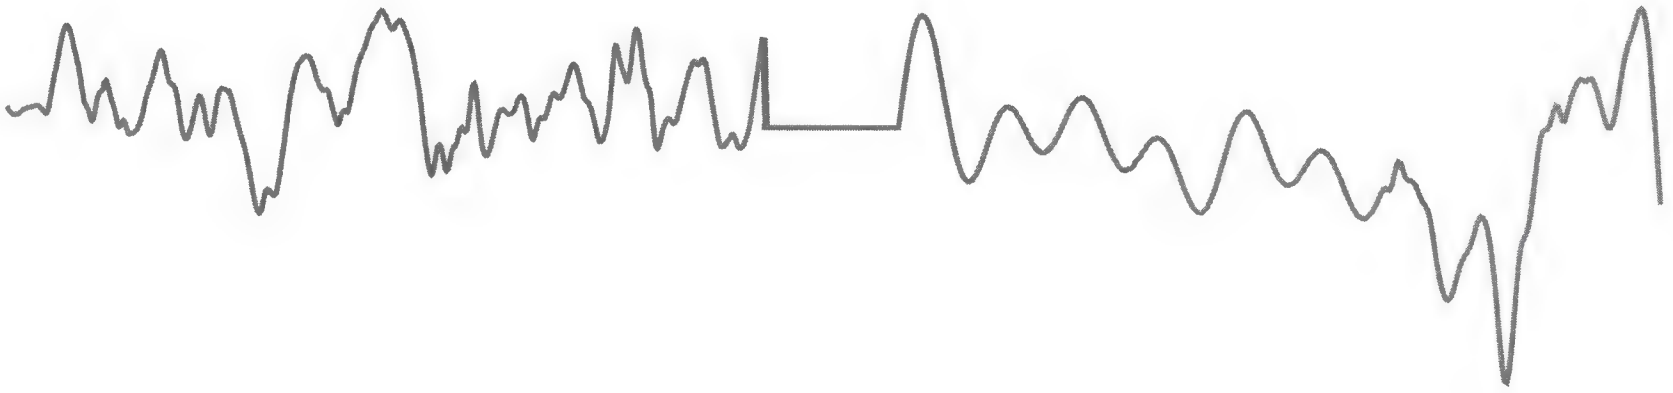

Aperiodic Example

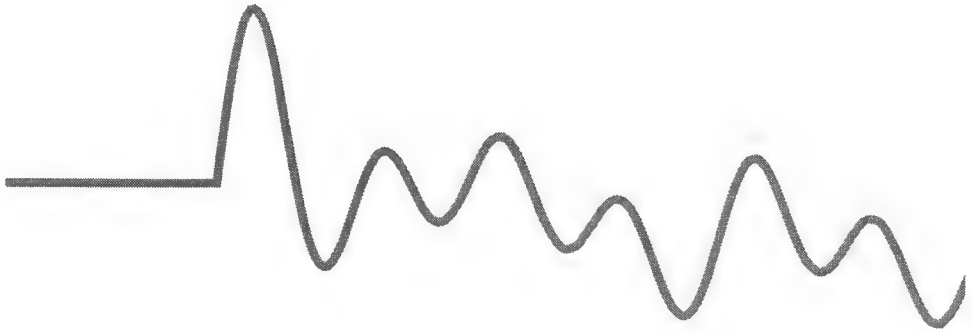

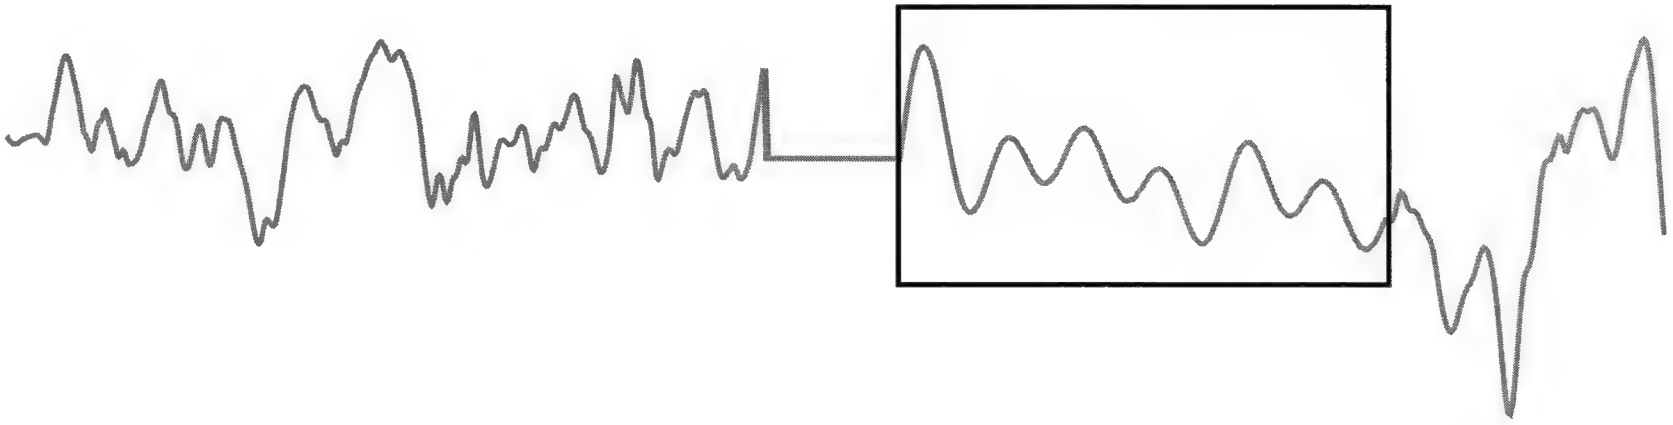

Aperiodic Example

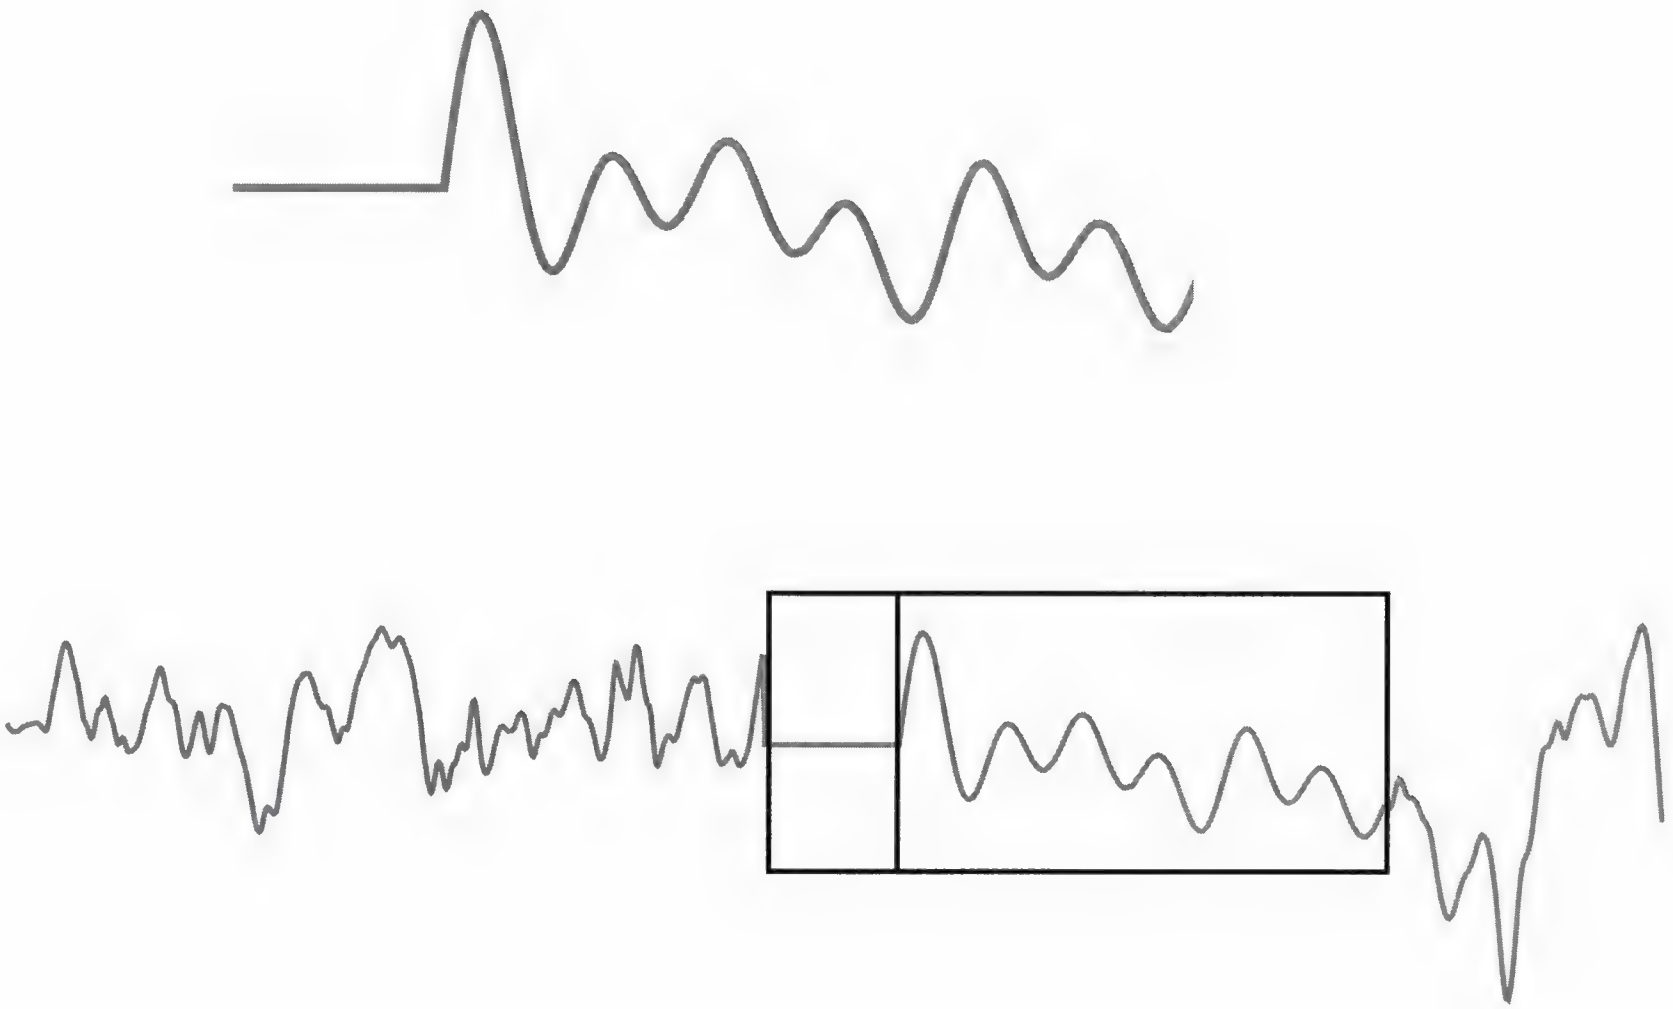

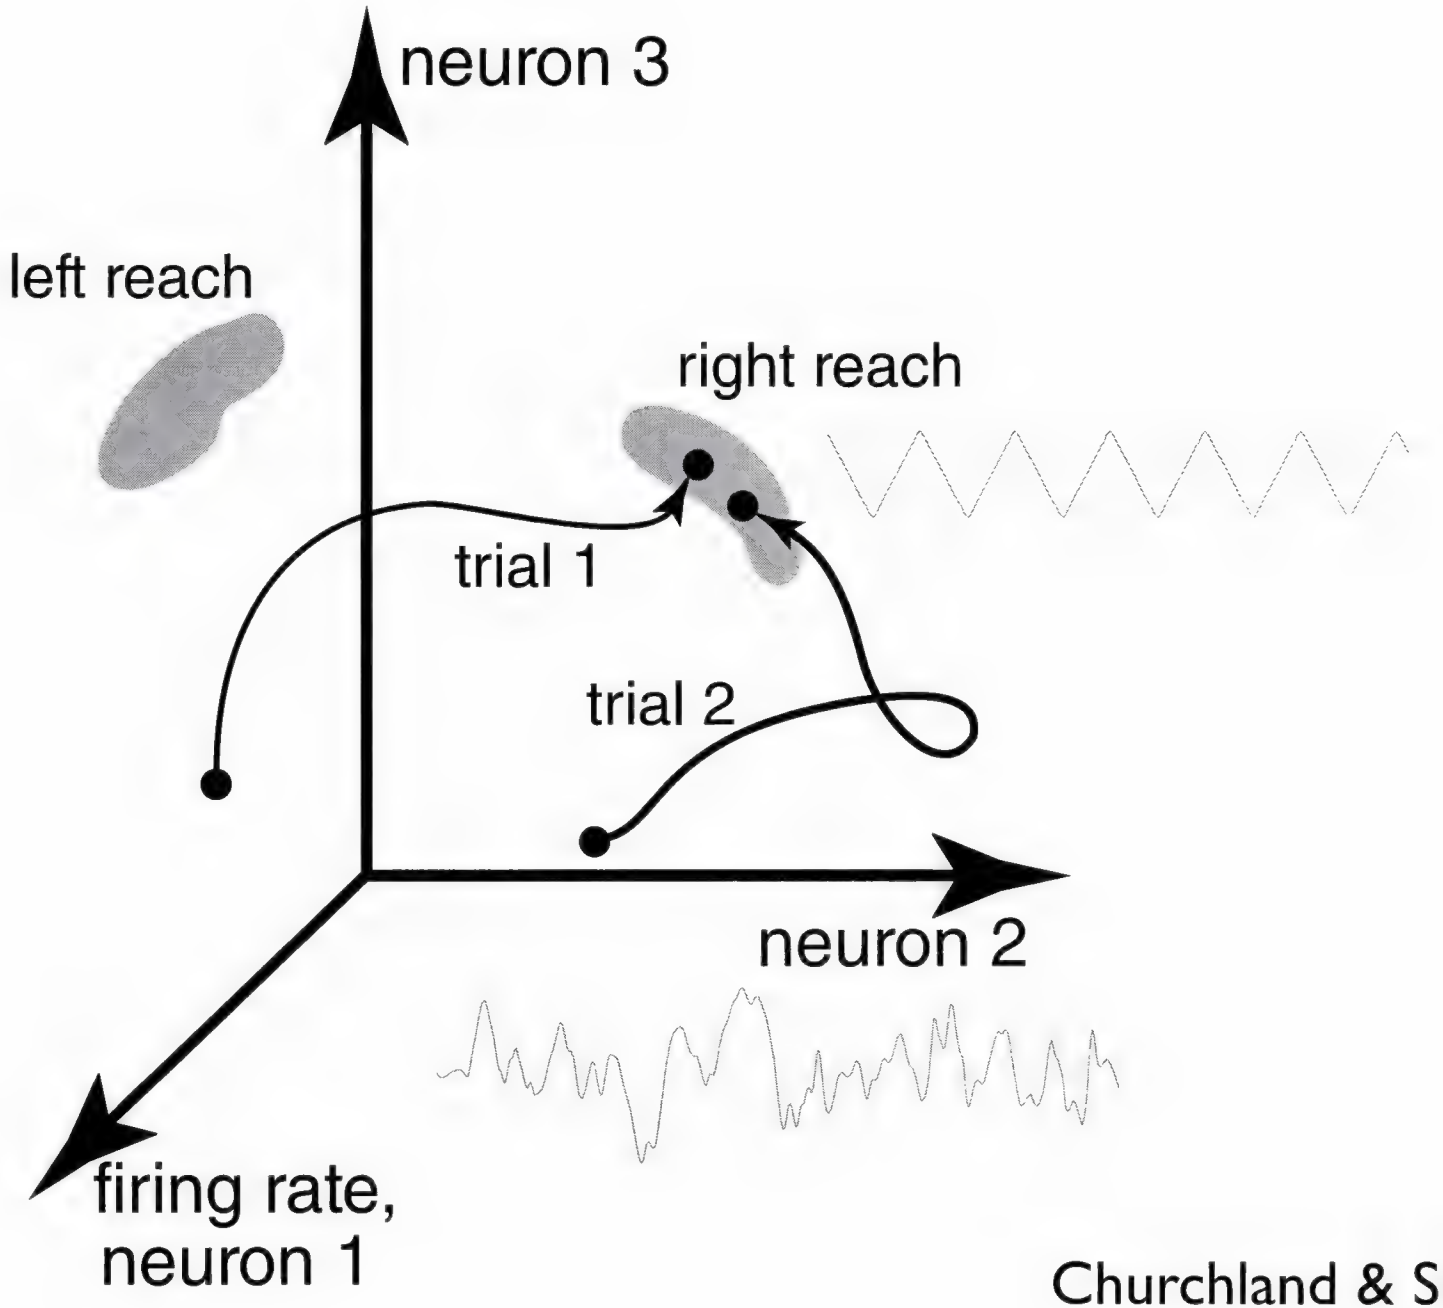

Churchland & Shenoy
J. Neuroscience 2006

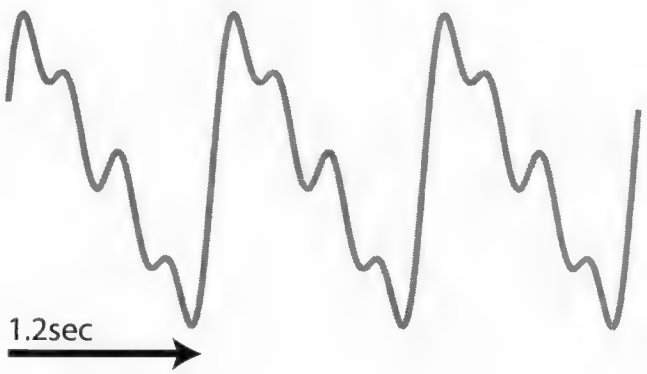
1.2sec

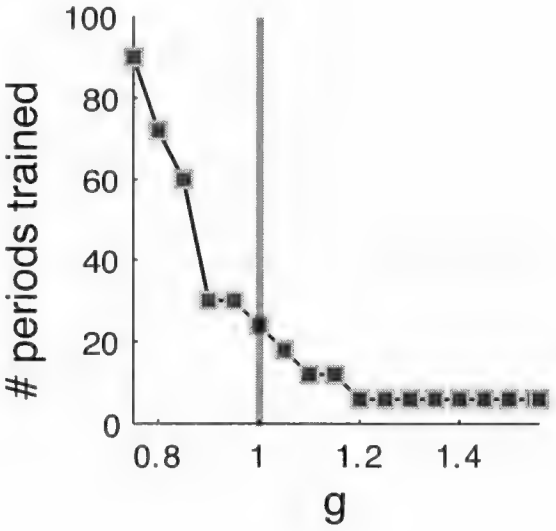
100
80
60
40
20
0
periods trained
0.8
1
1.2
1.4
g

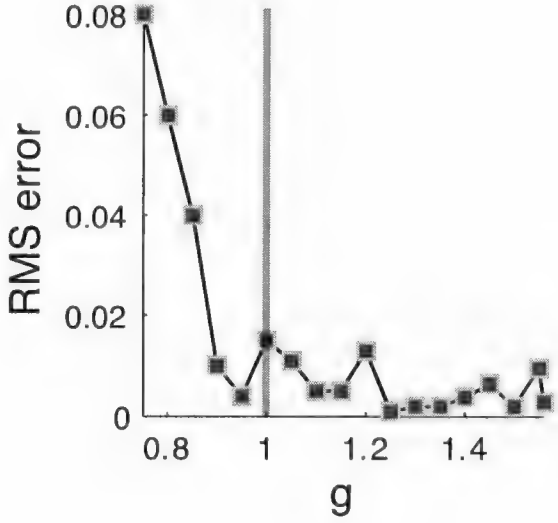
0.08
0.06
0.04
0.02
0
RMS error
0.8
1
1.2
1.4
g

$|w_o|$

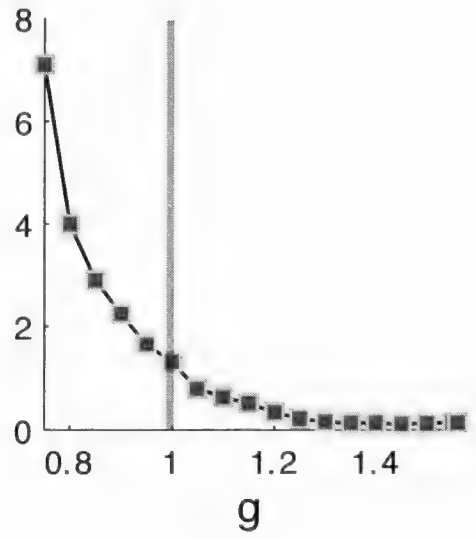
8
6
4
2
0
0.8
1
1.2
1.4
g

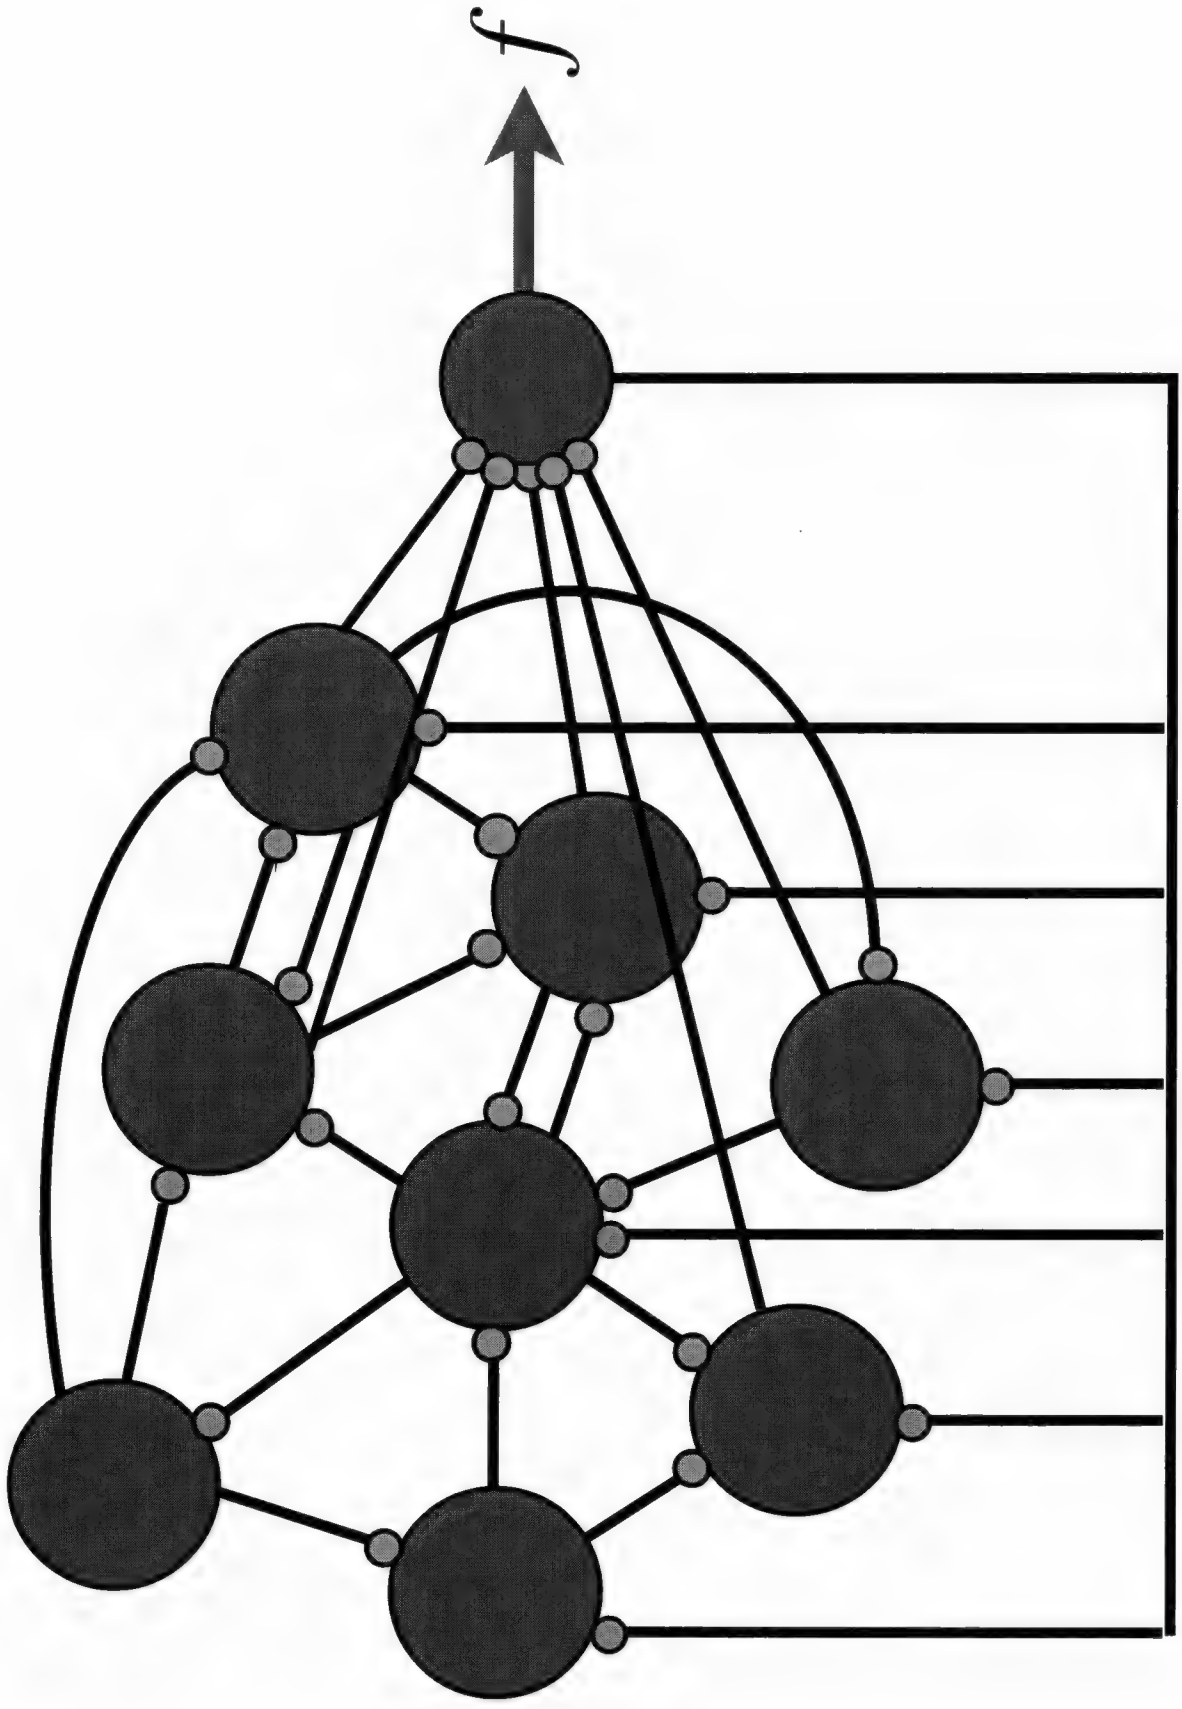

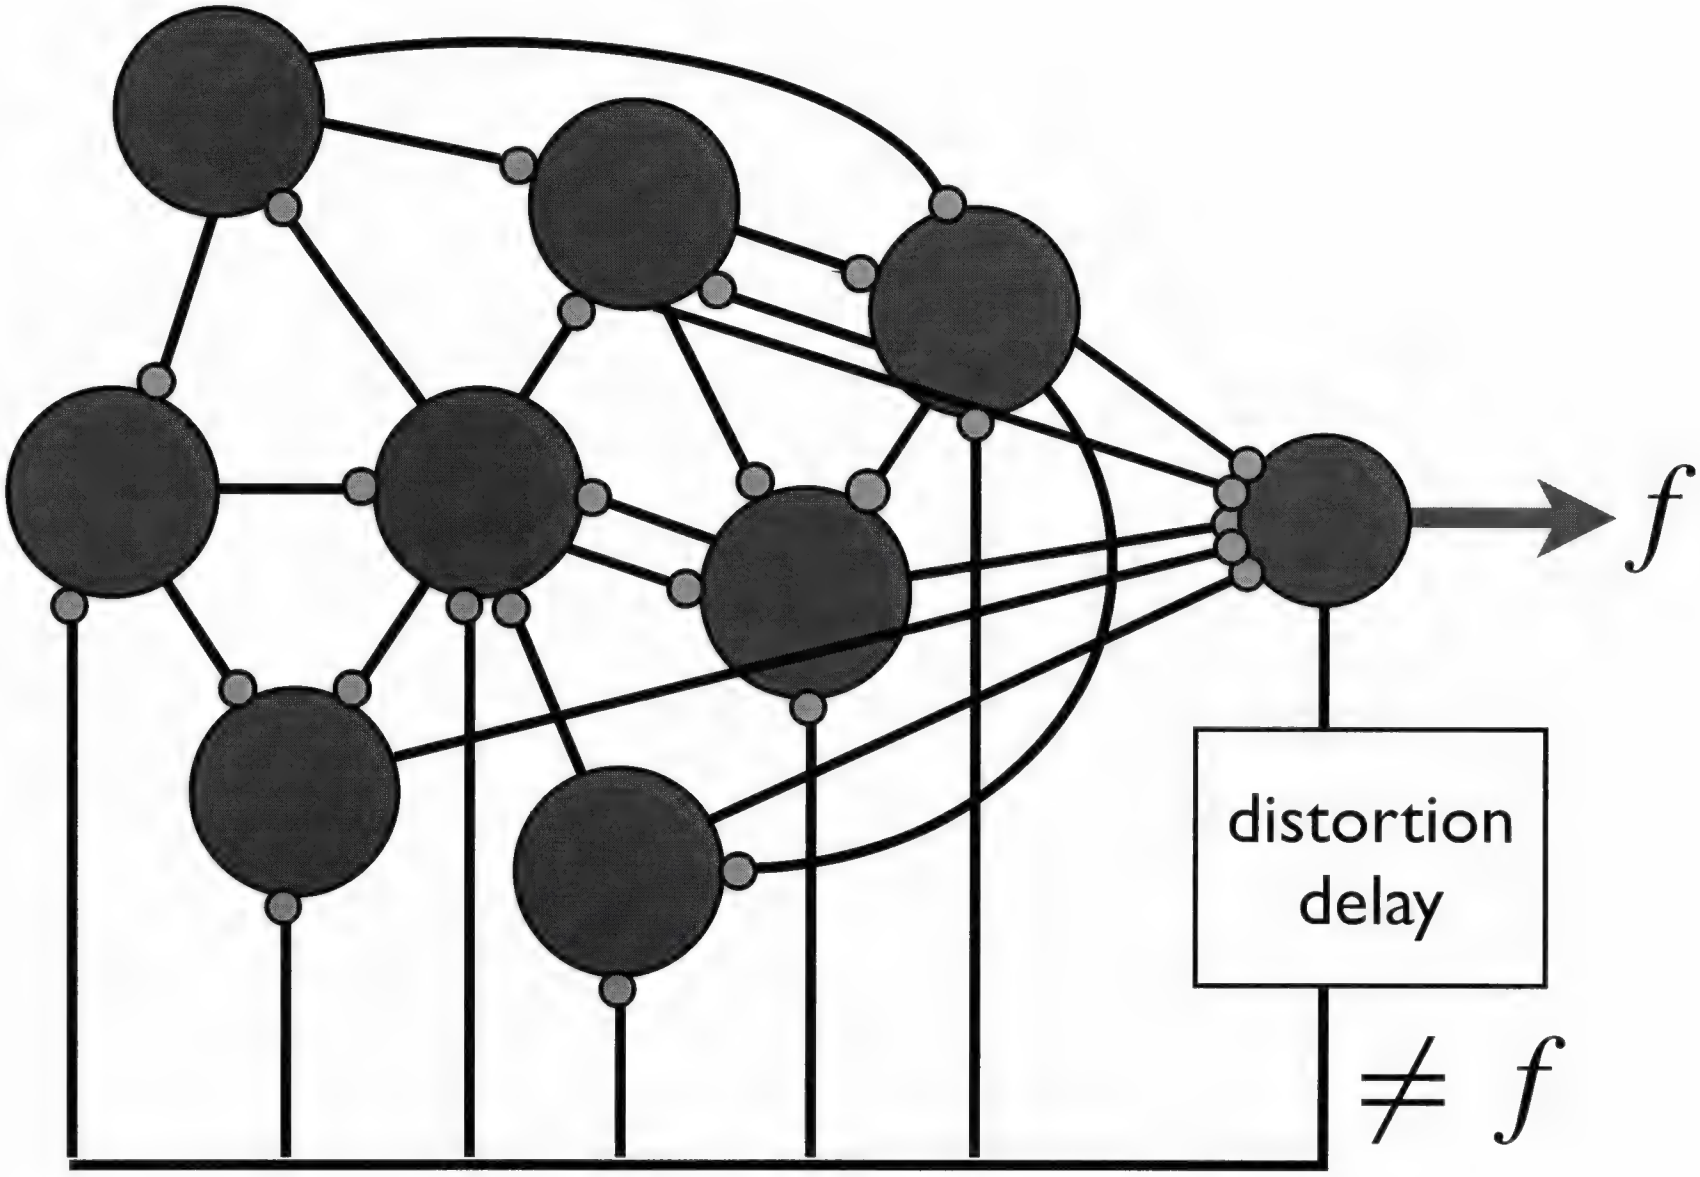
f
distortion
delay
$\neq f$

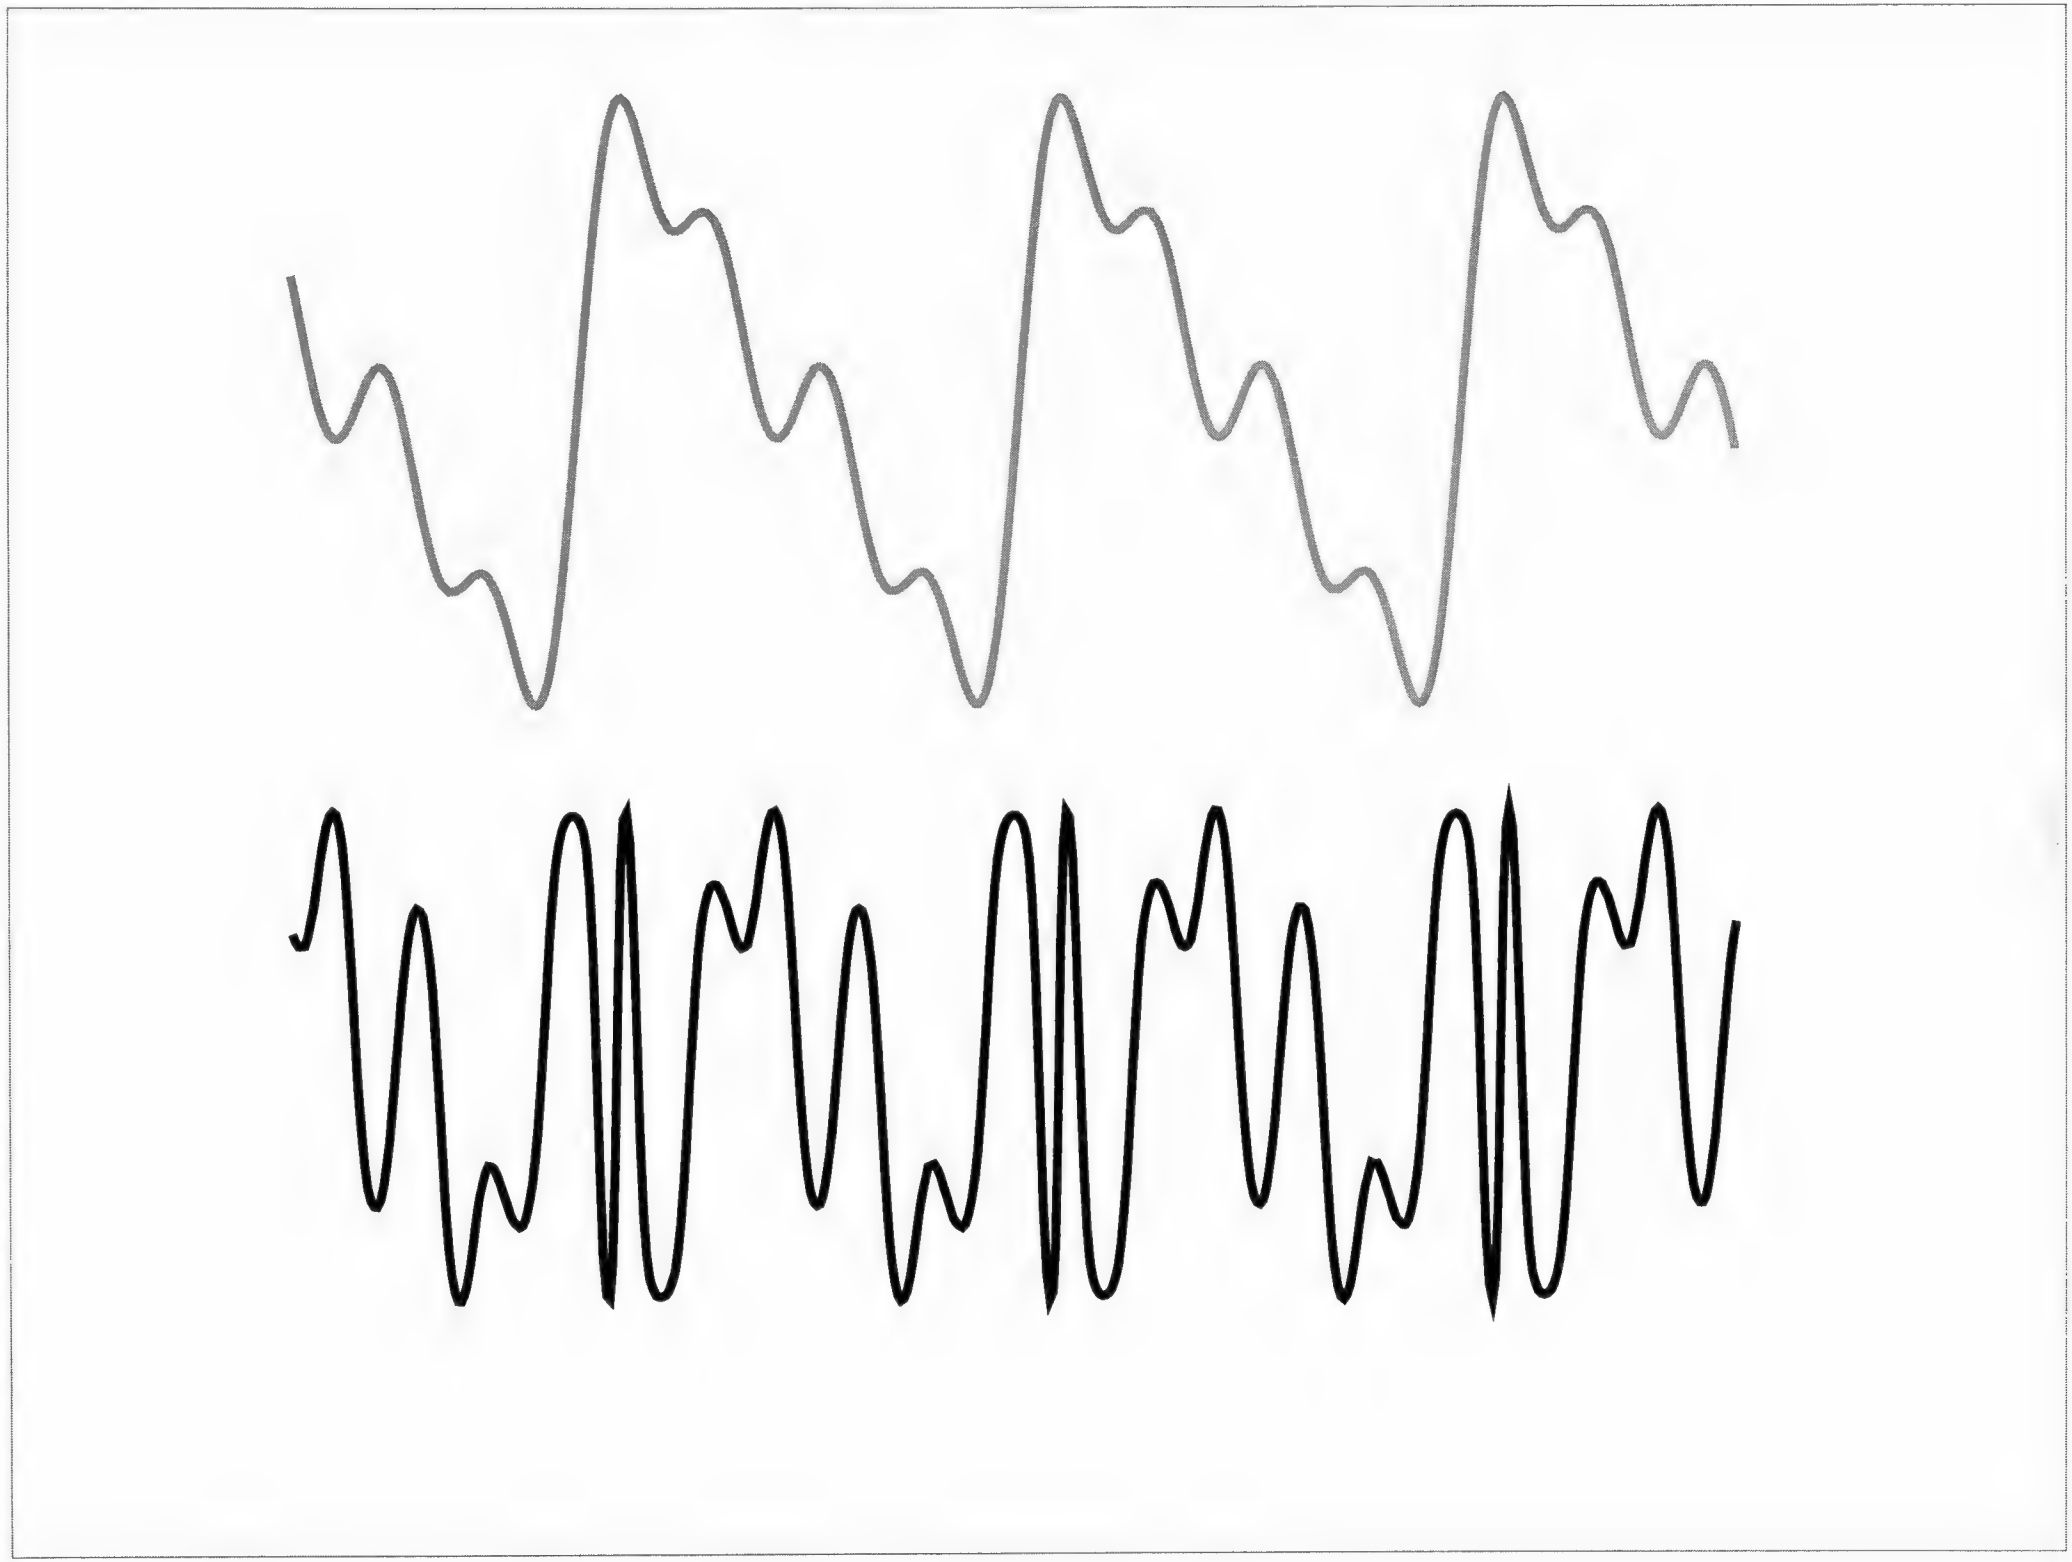

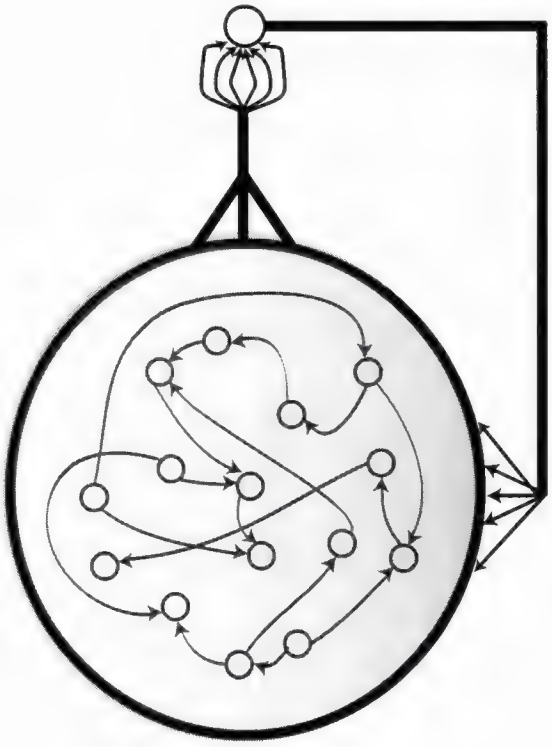

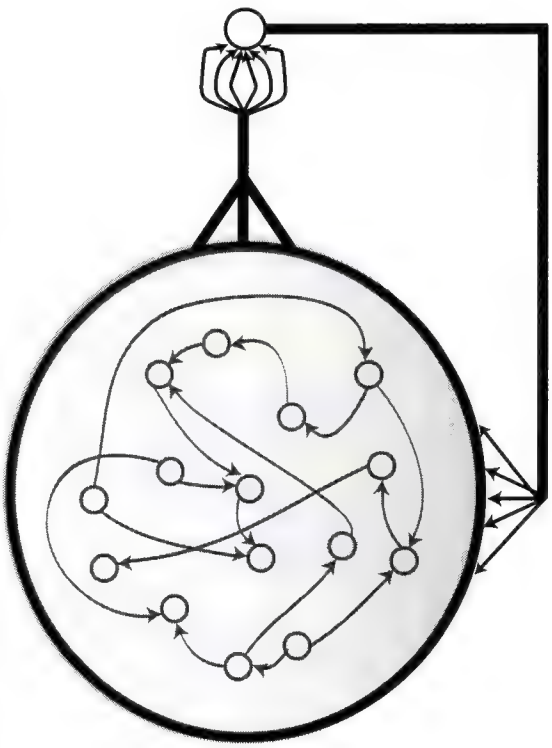

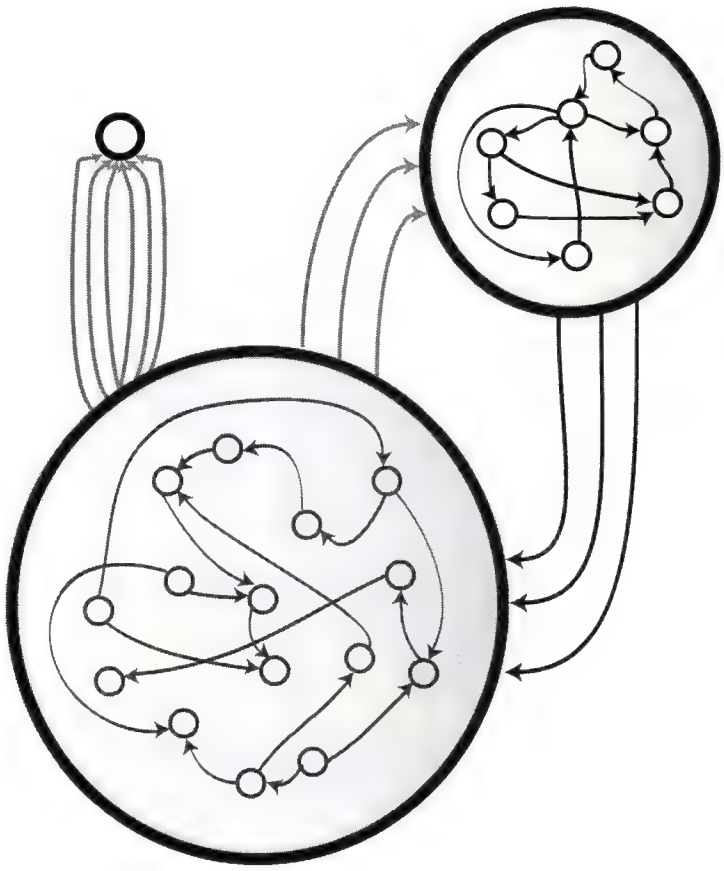

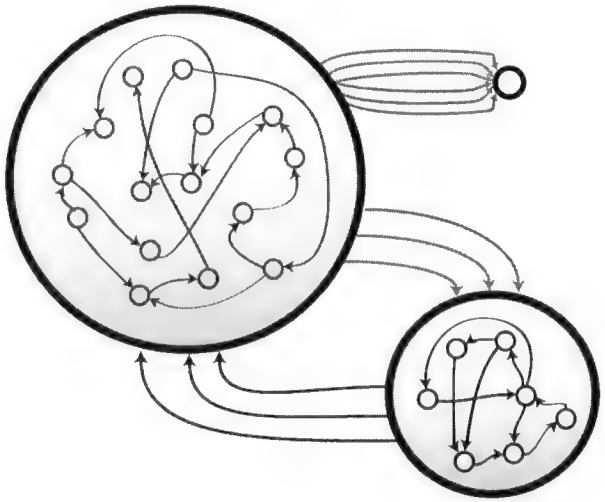

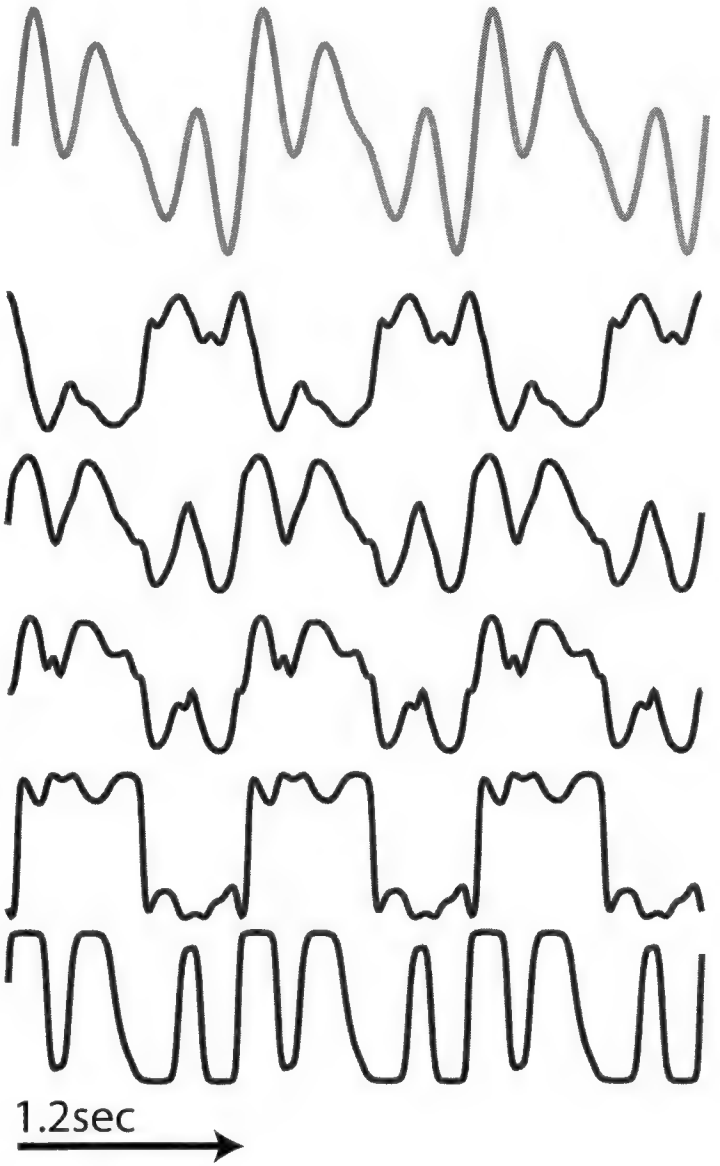
1.2sec

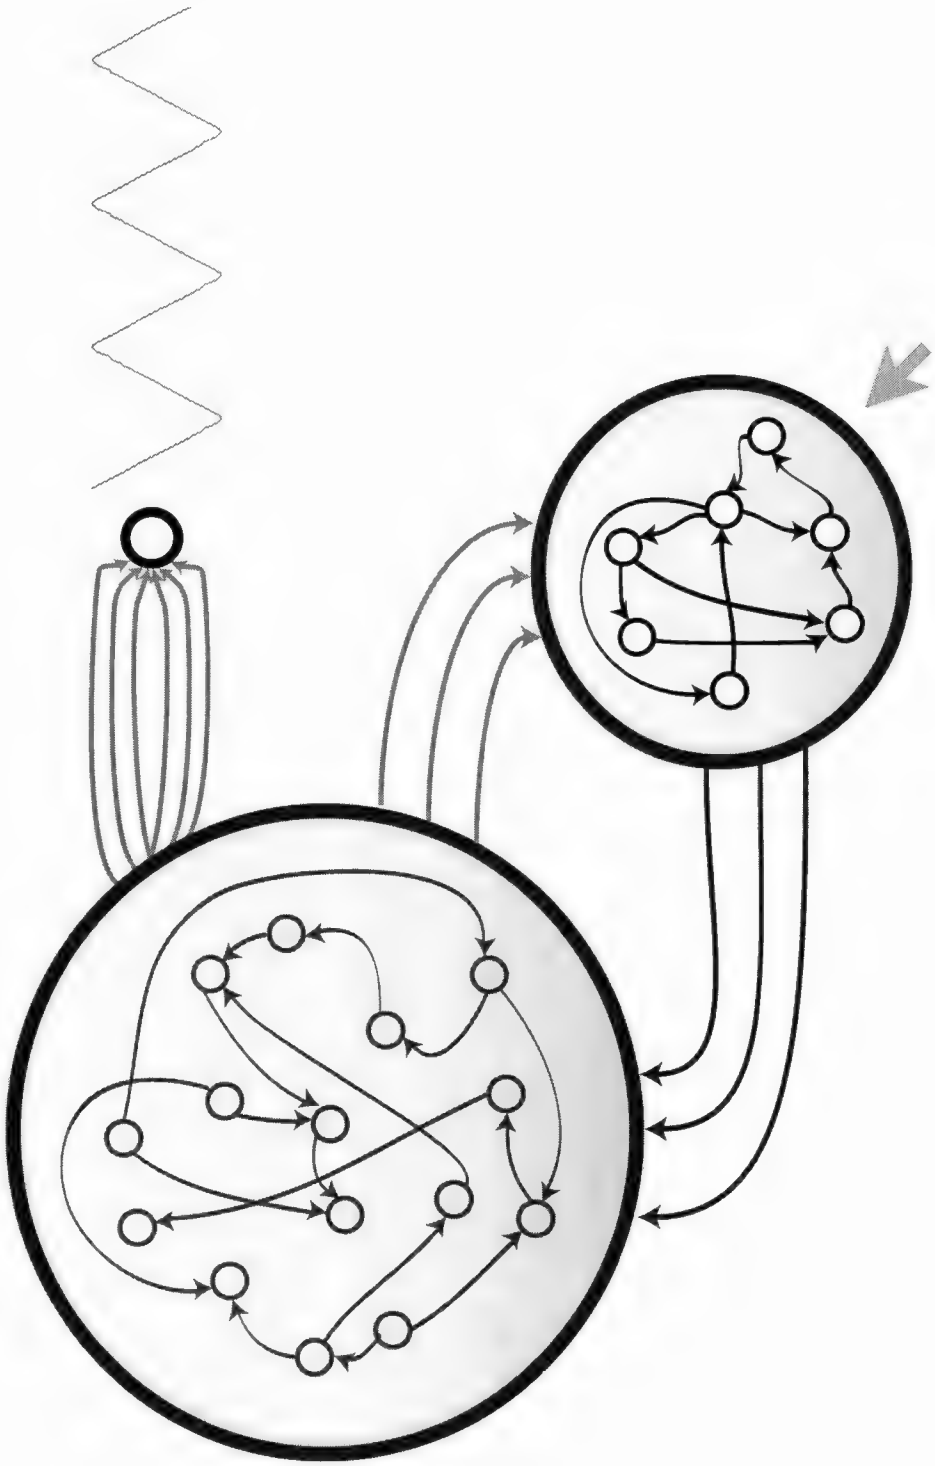

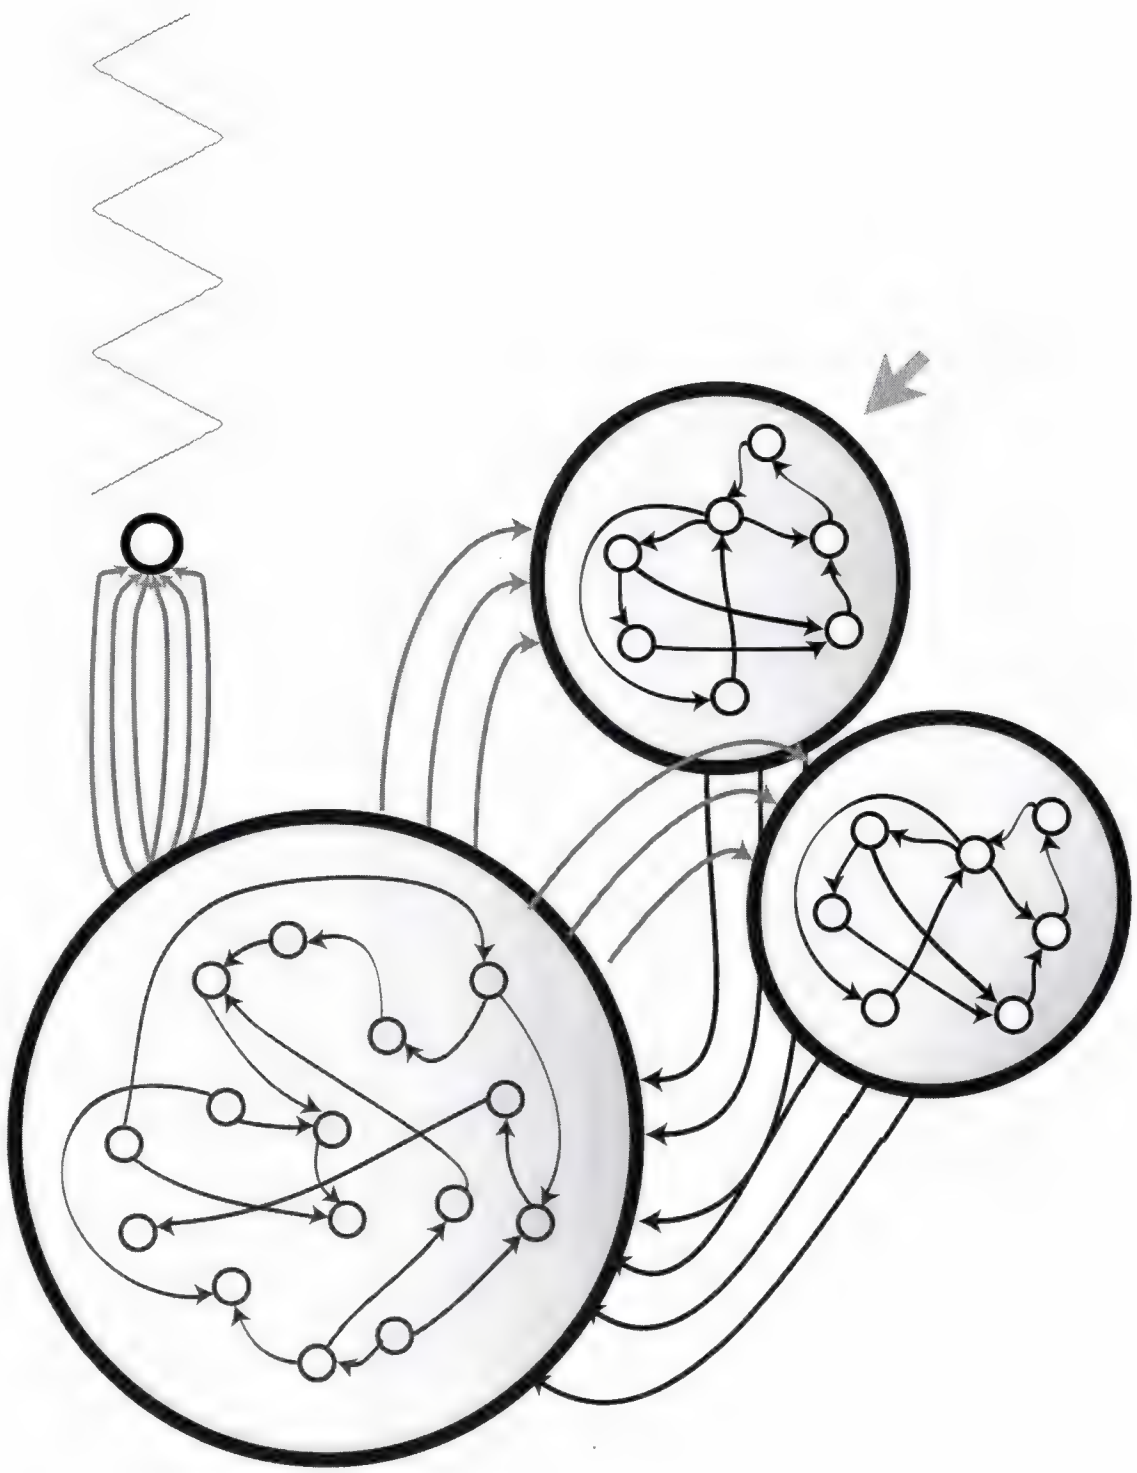

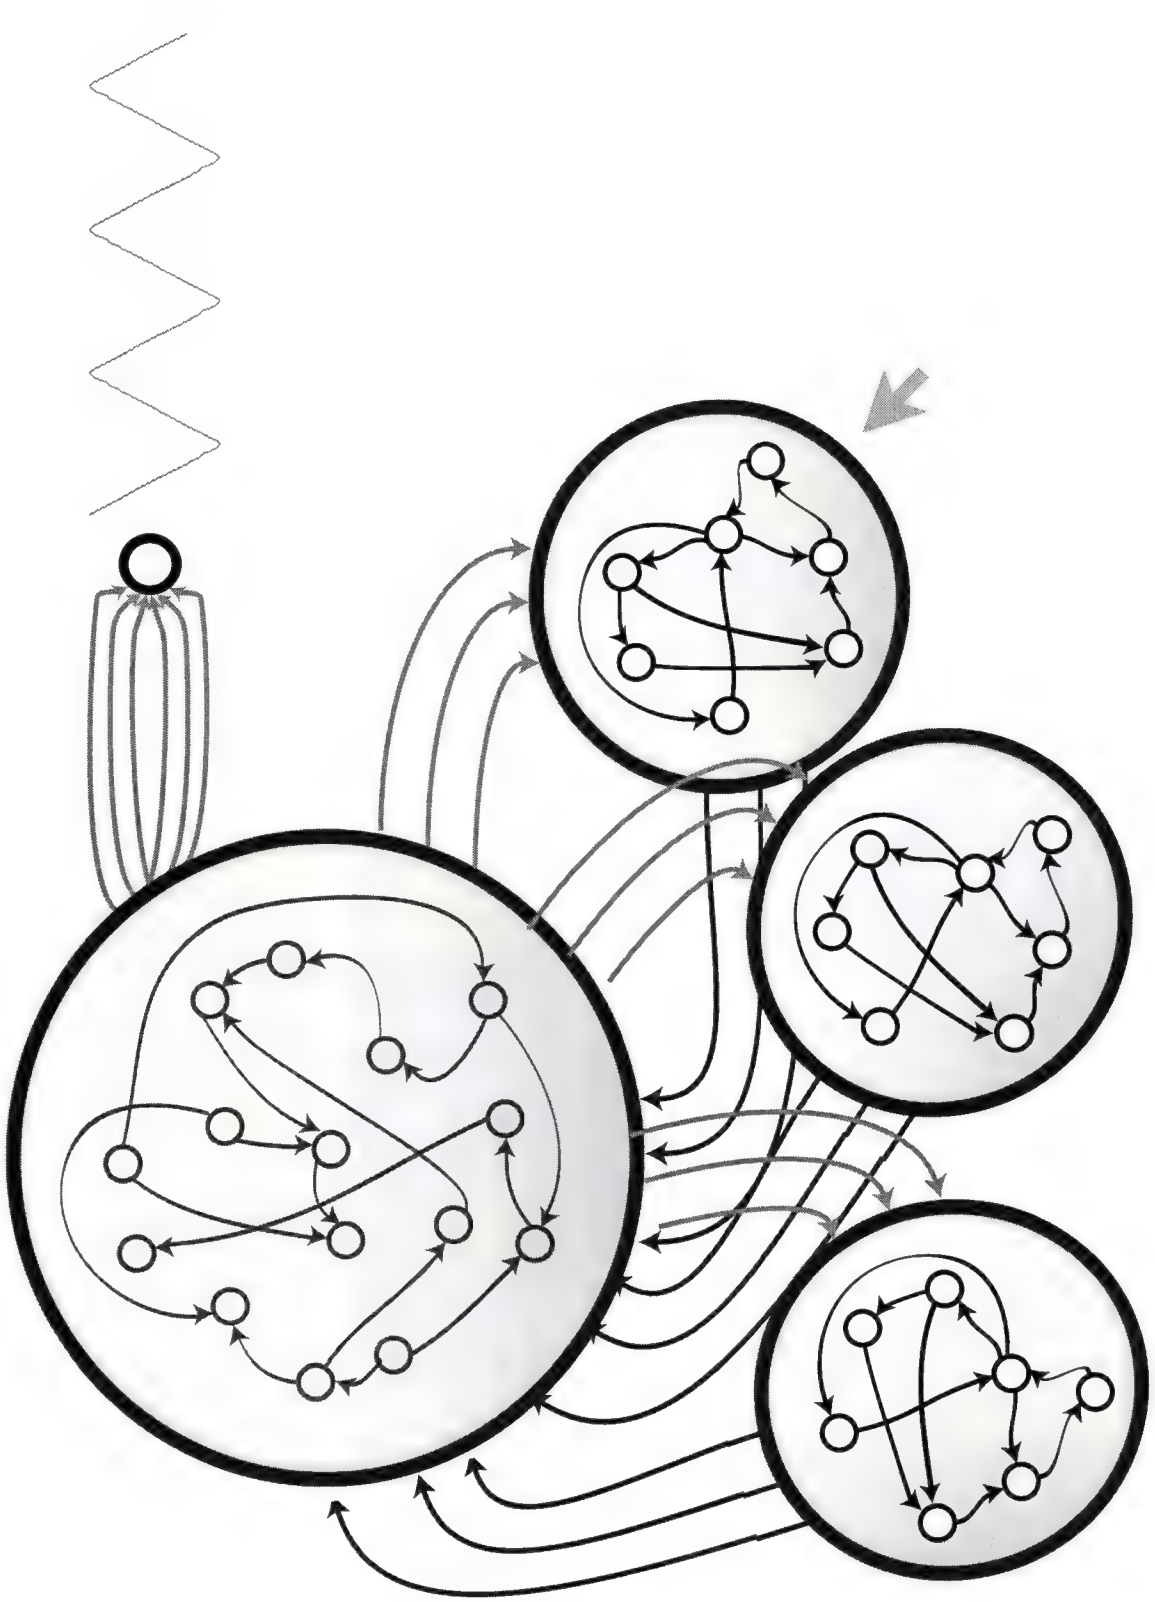

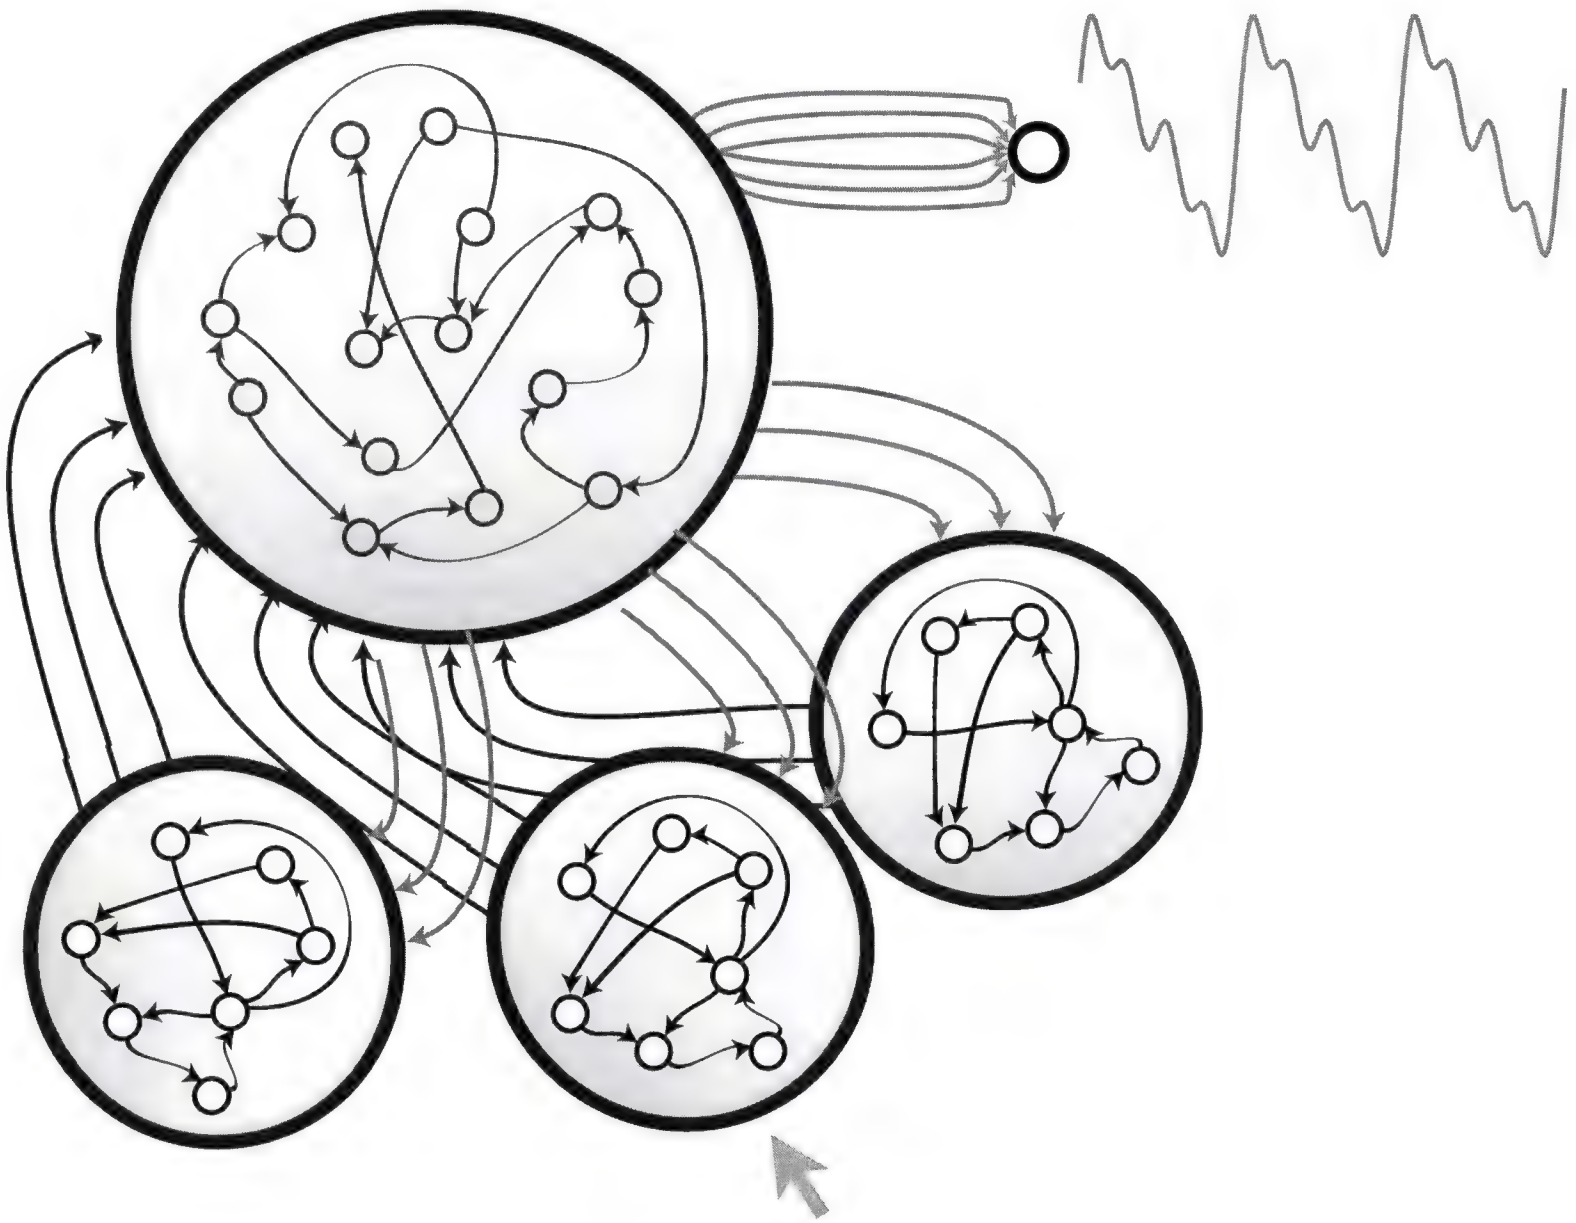

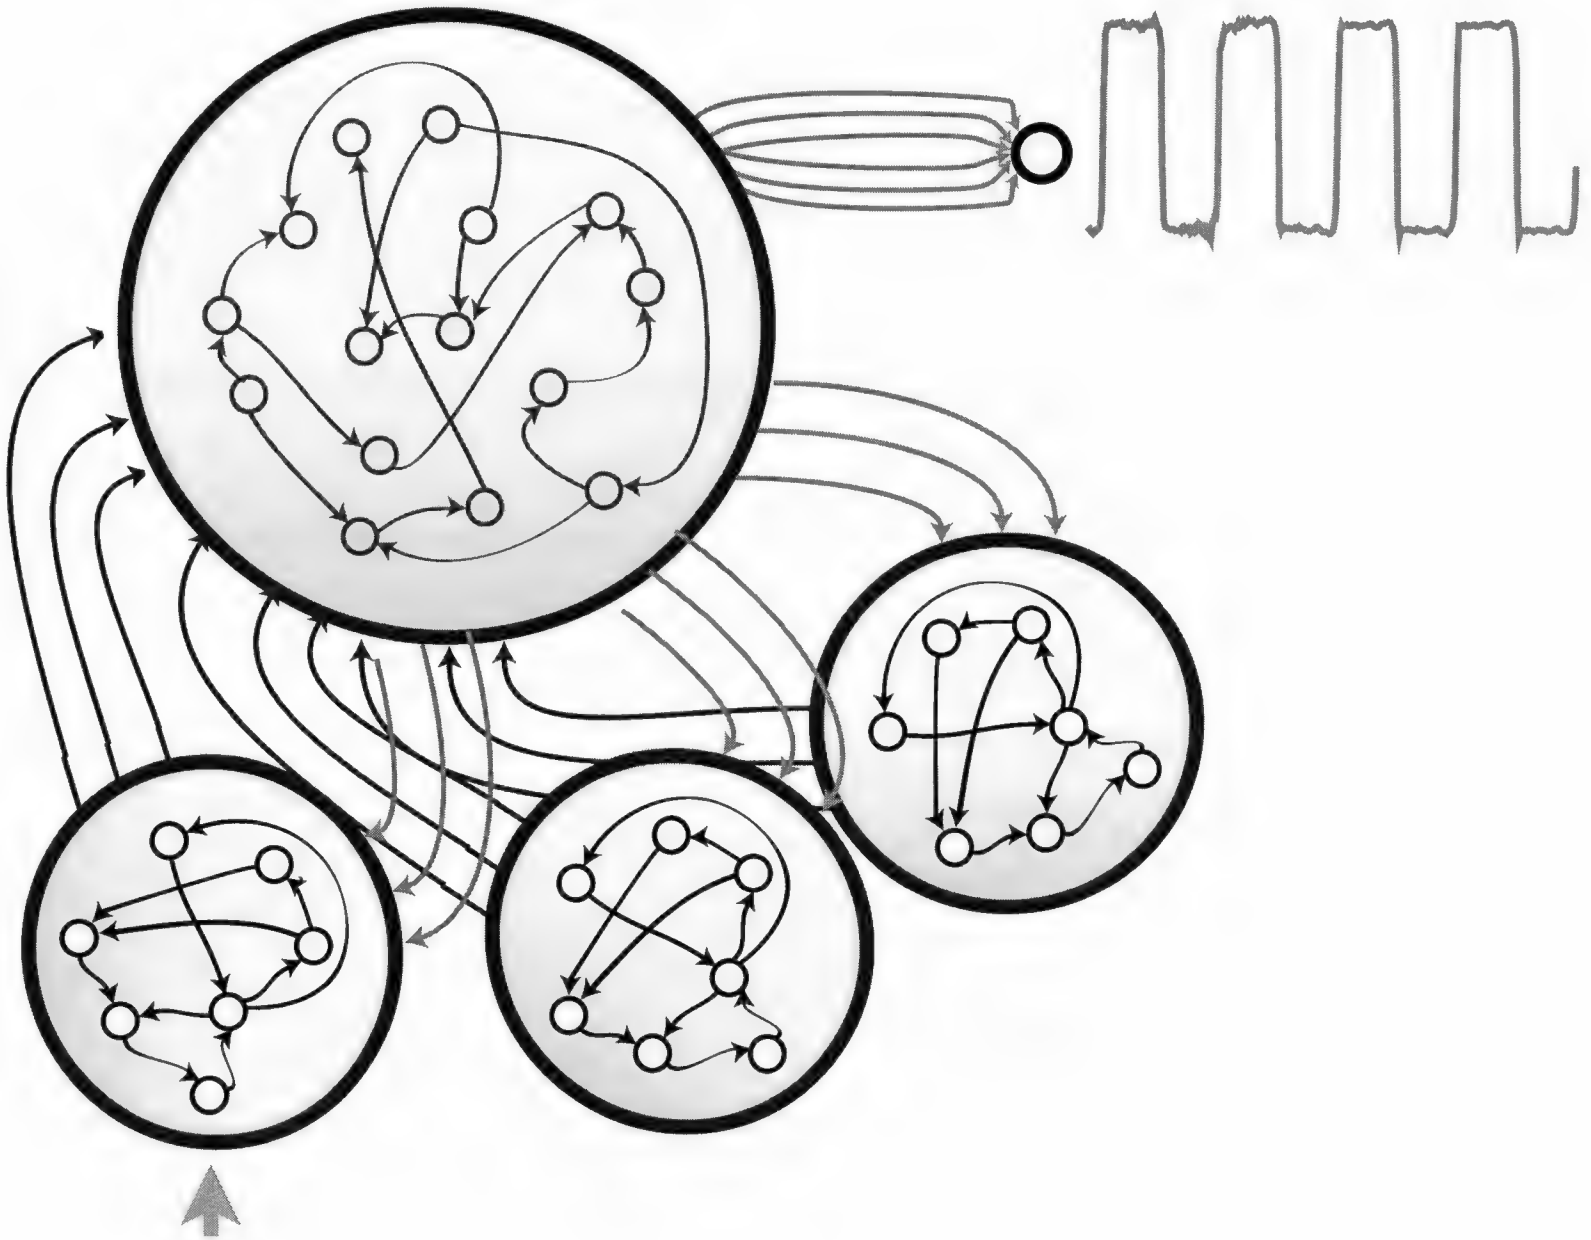

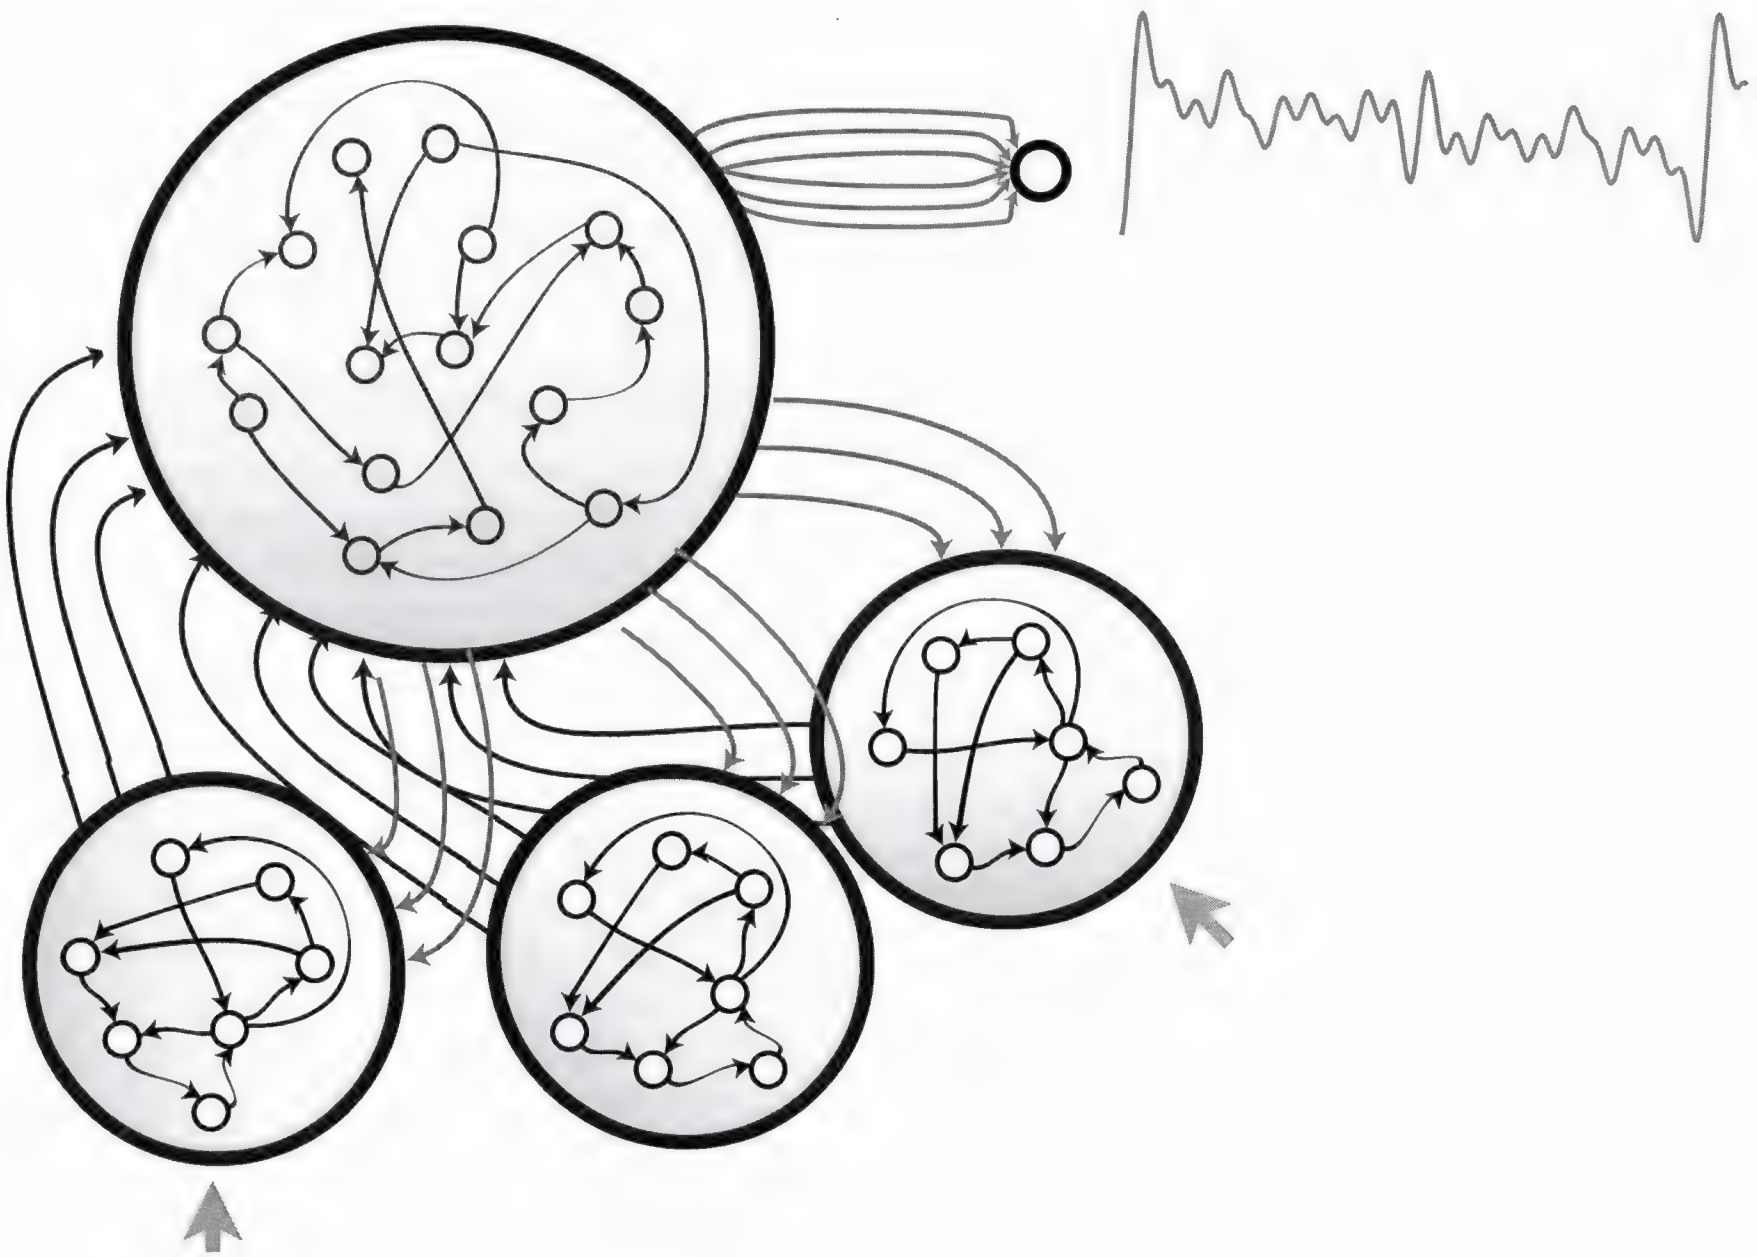

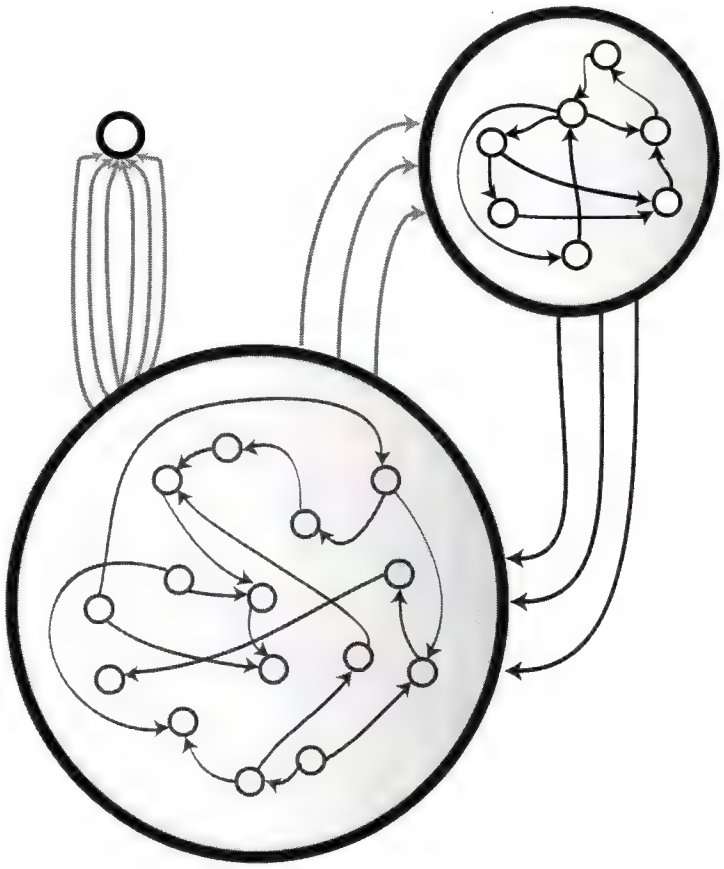

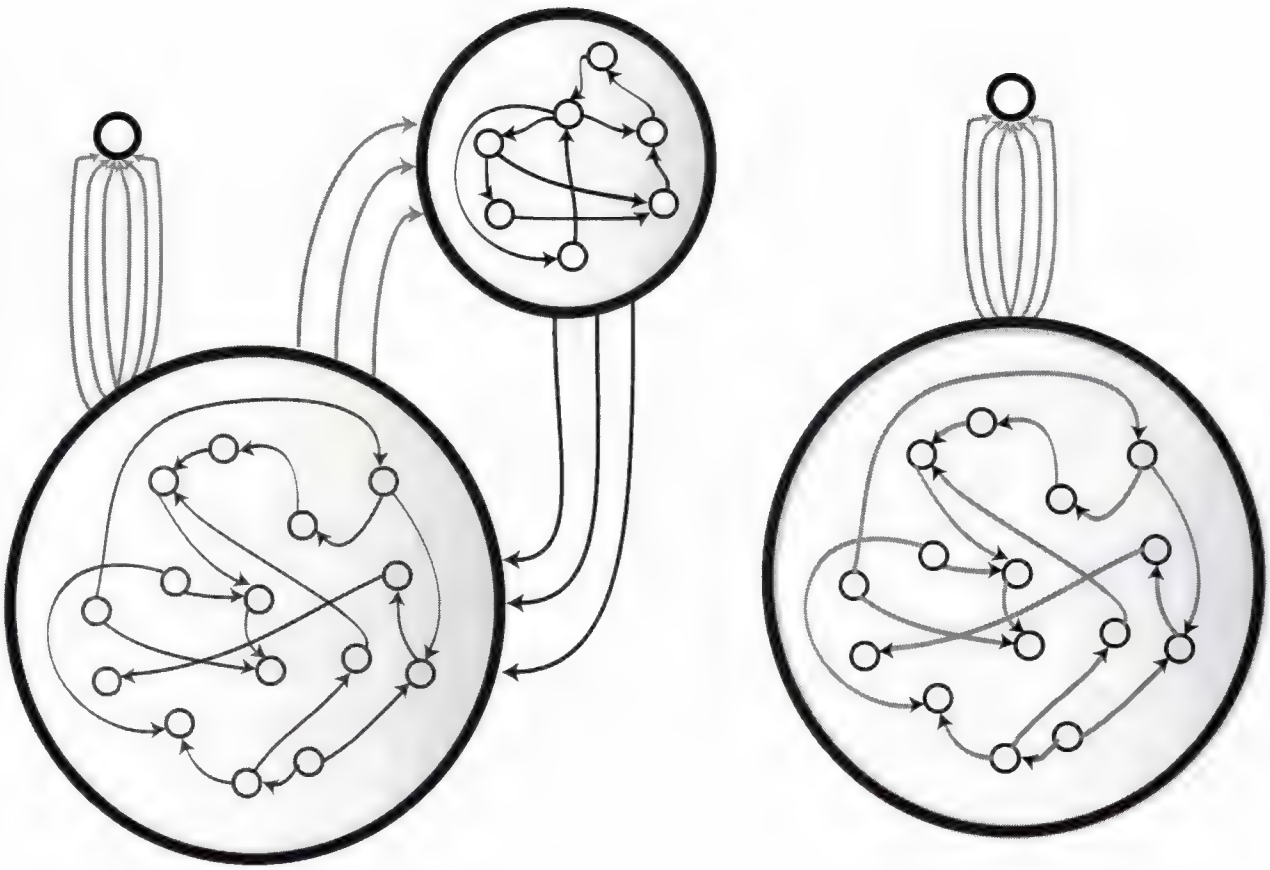

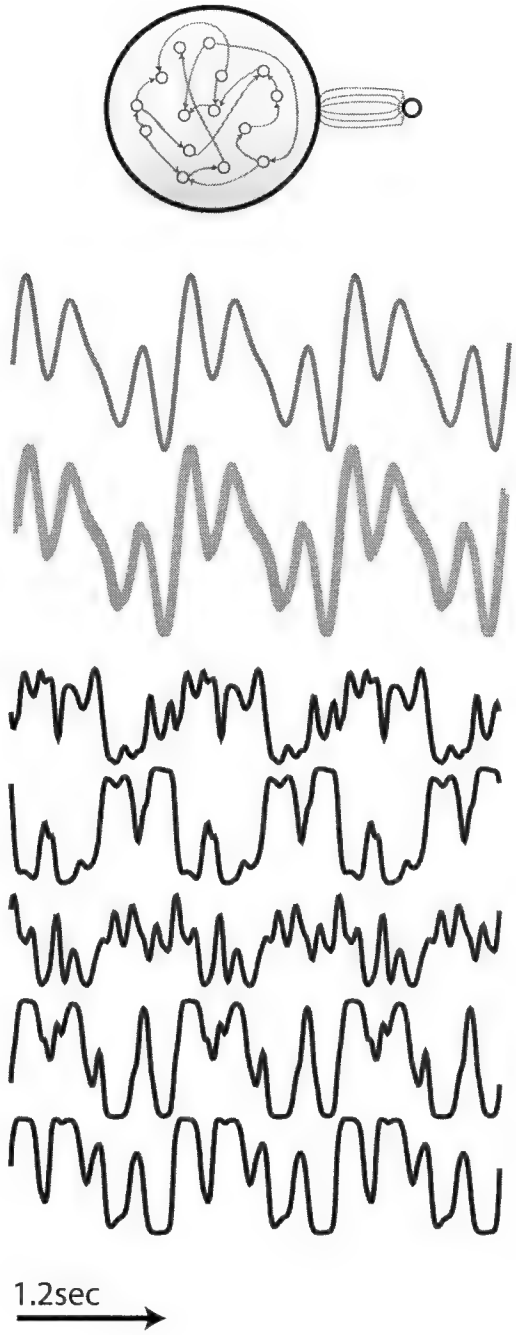
1.2sec

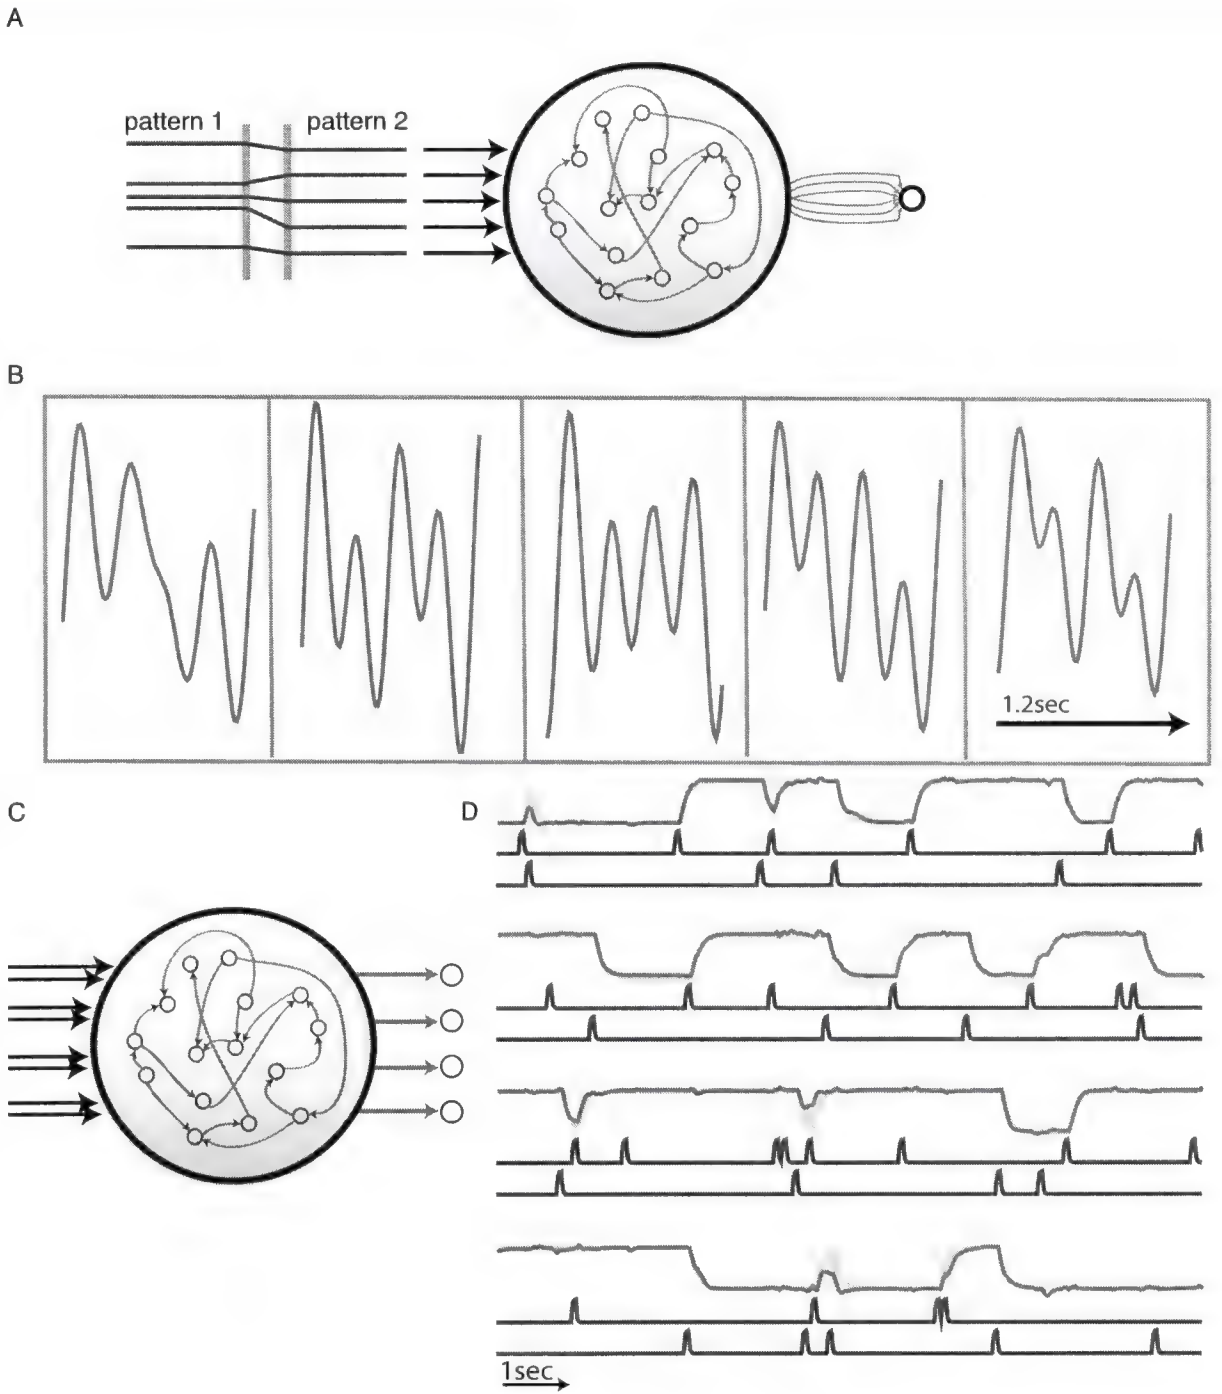
A
pattern 1
pattern 2
B
1.2sec
C
D
1sec

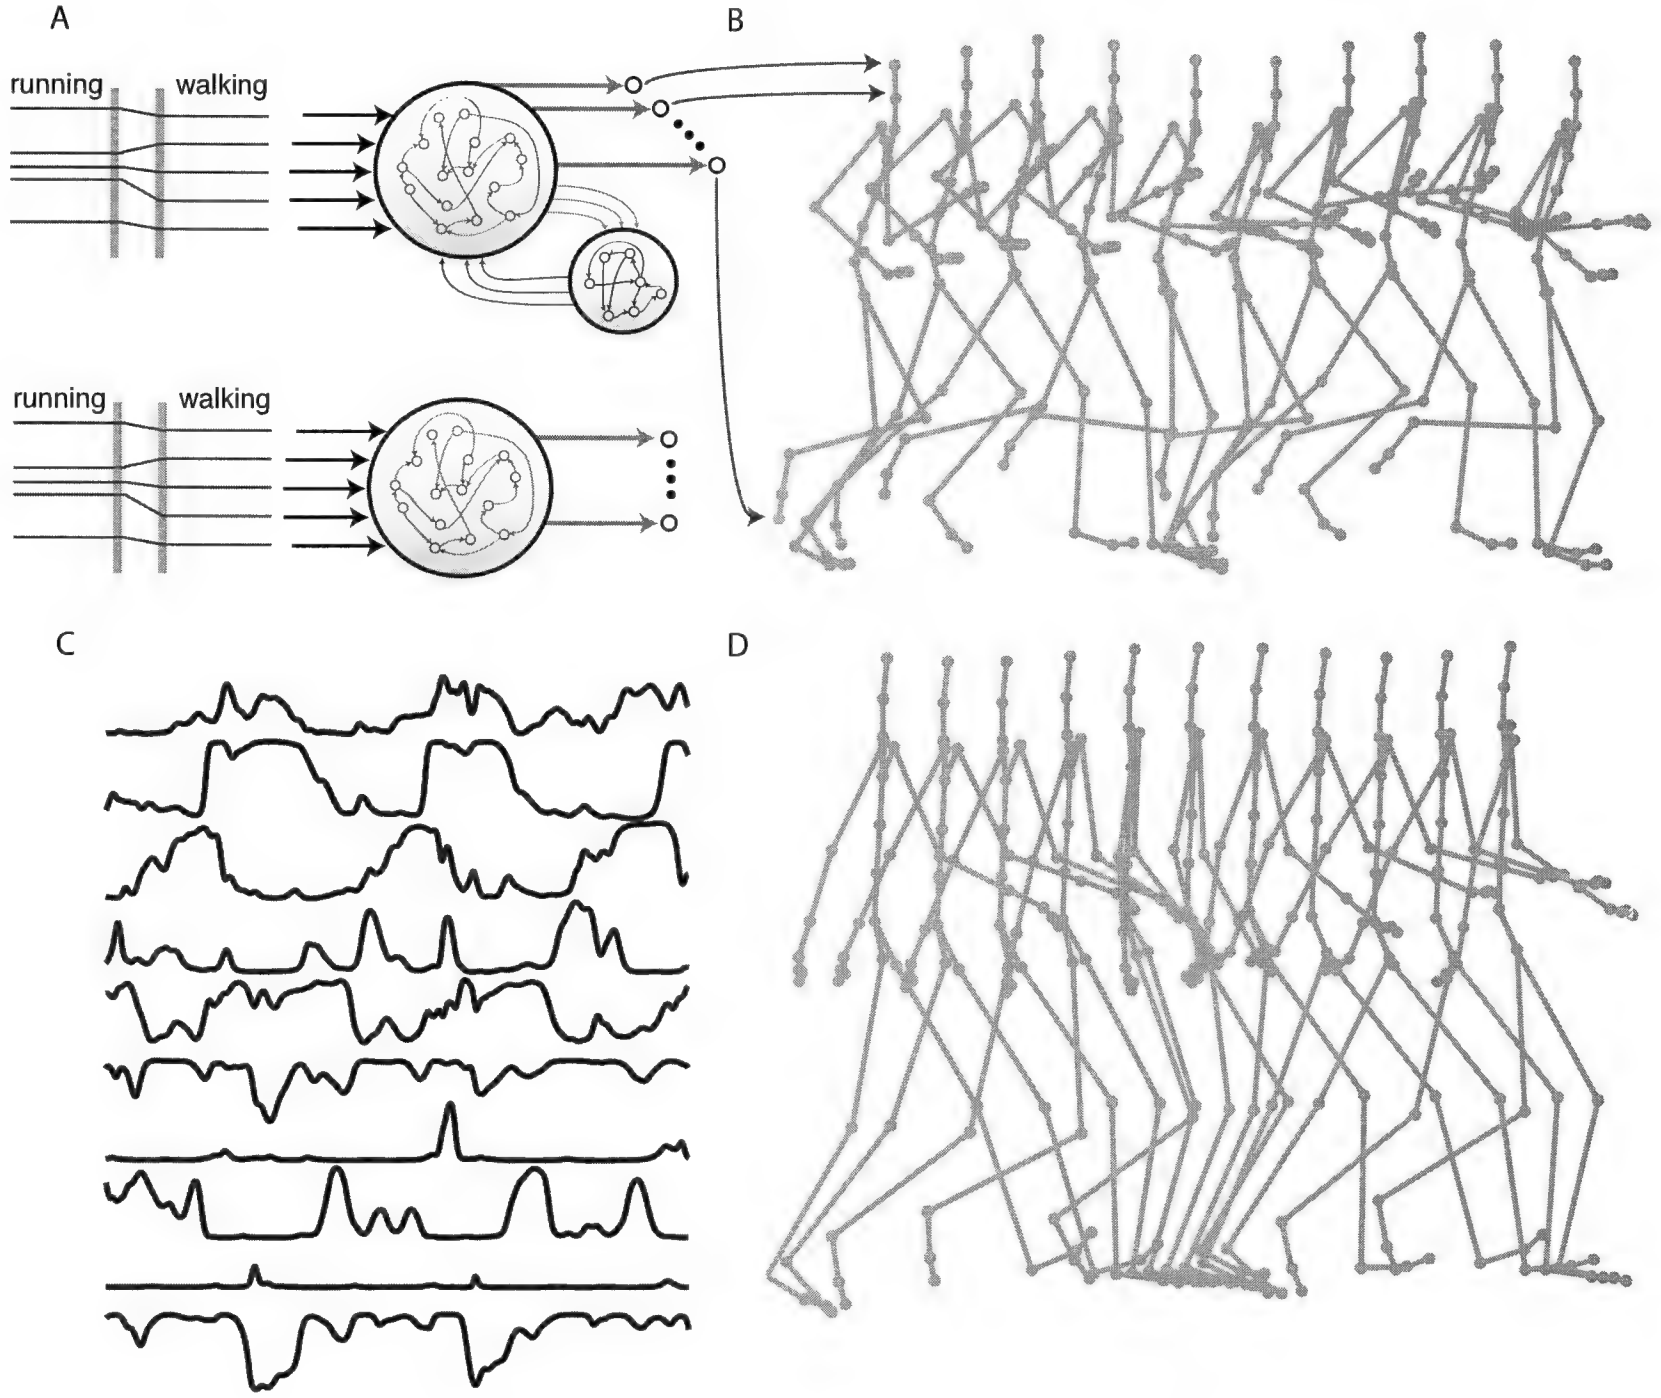
A
running
walking
running
walking
B
C
D

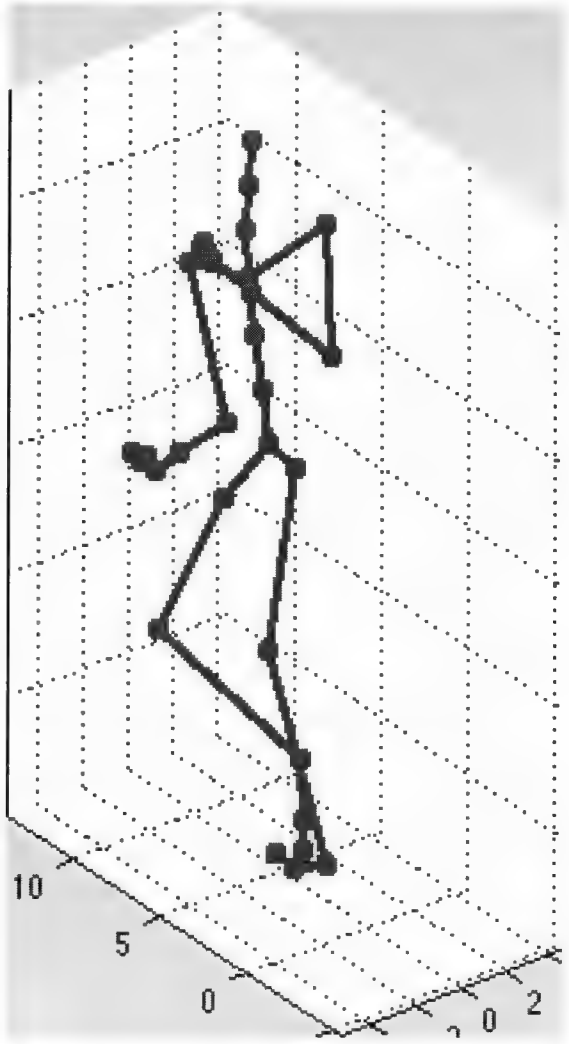
10
5
0
0
2

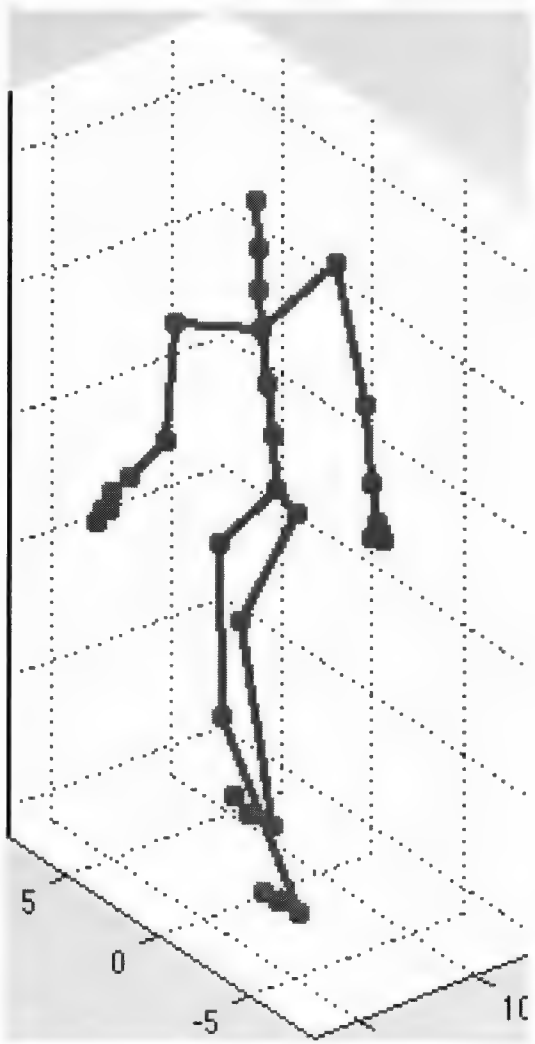
5
0
-5

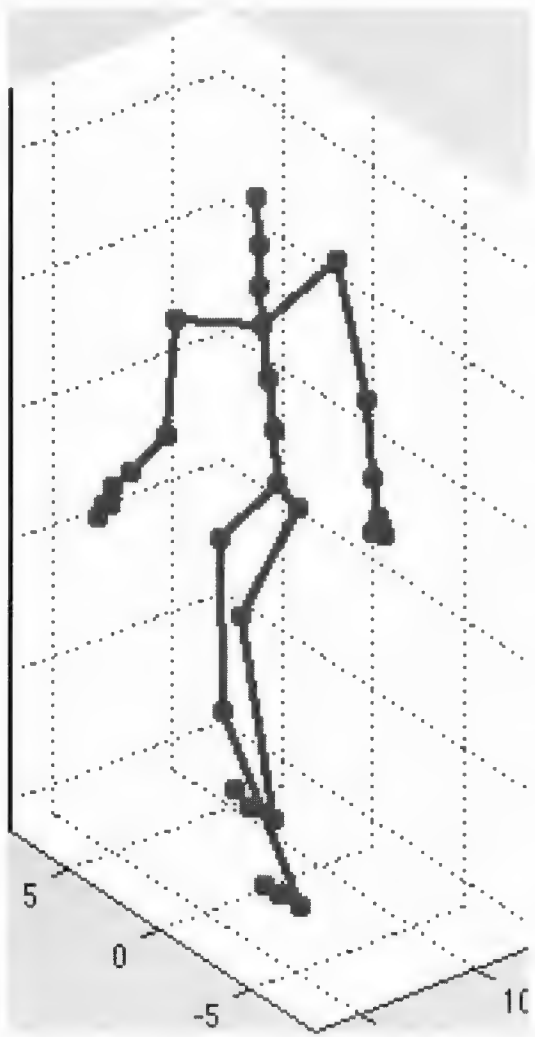
5
0
-5
10

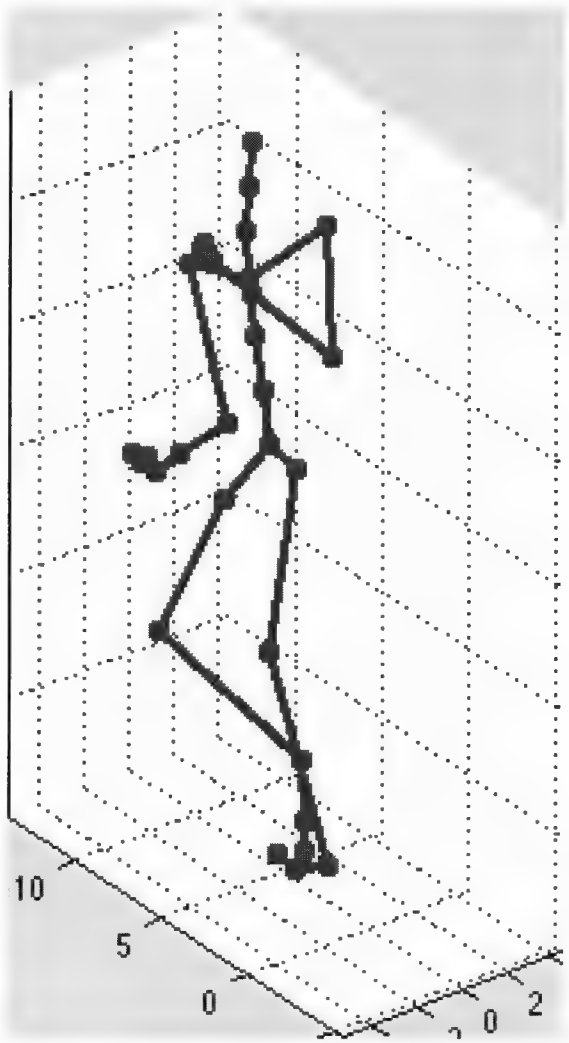
10
5
0
0
2

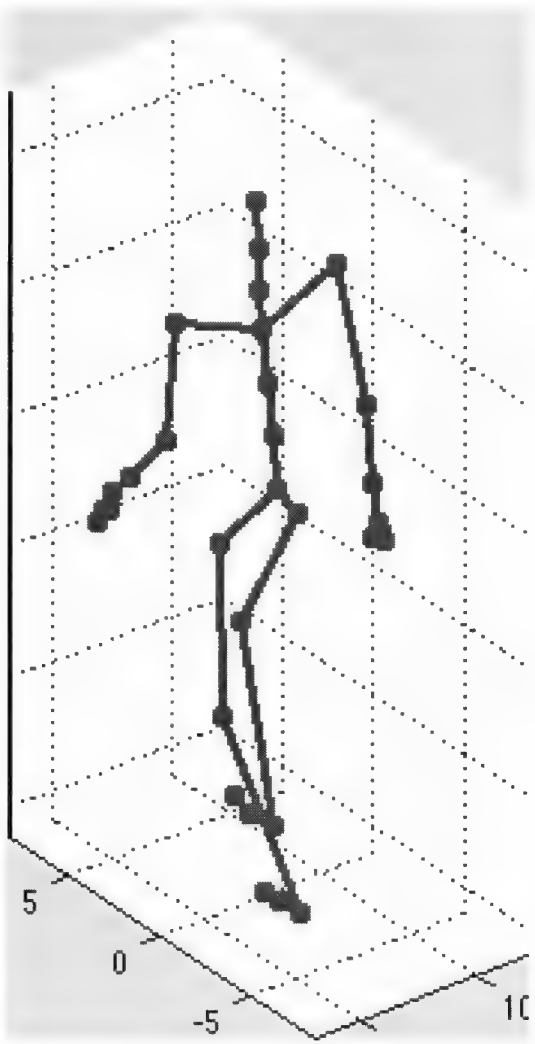
5
0
-5

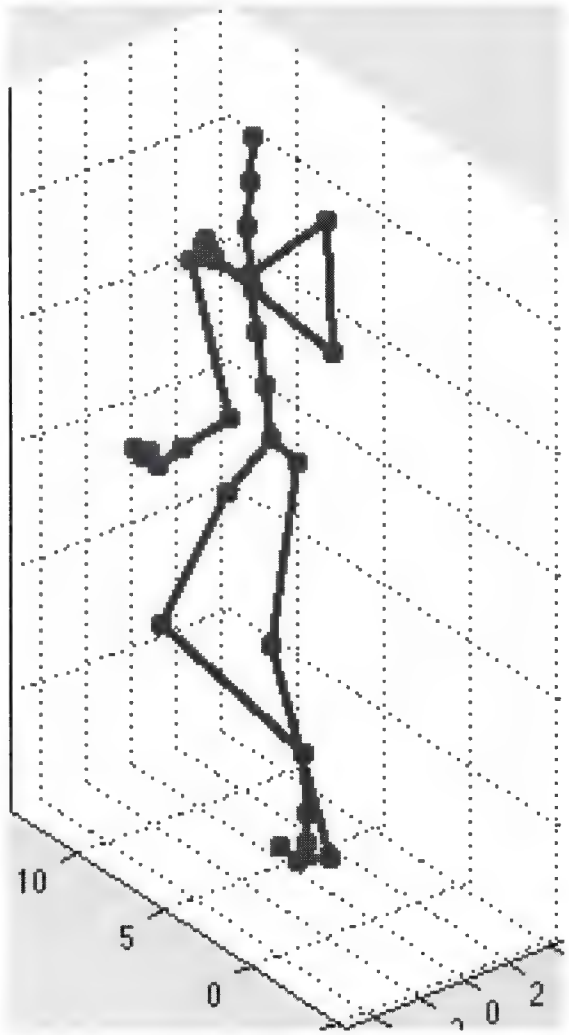
10
5
0
0
2

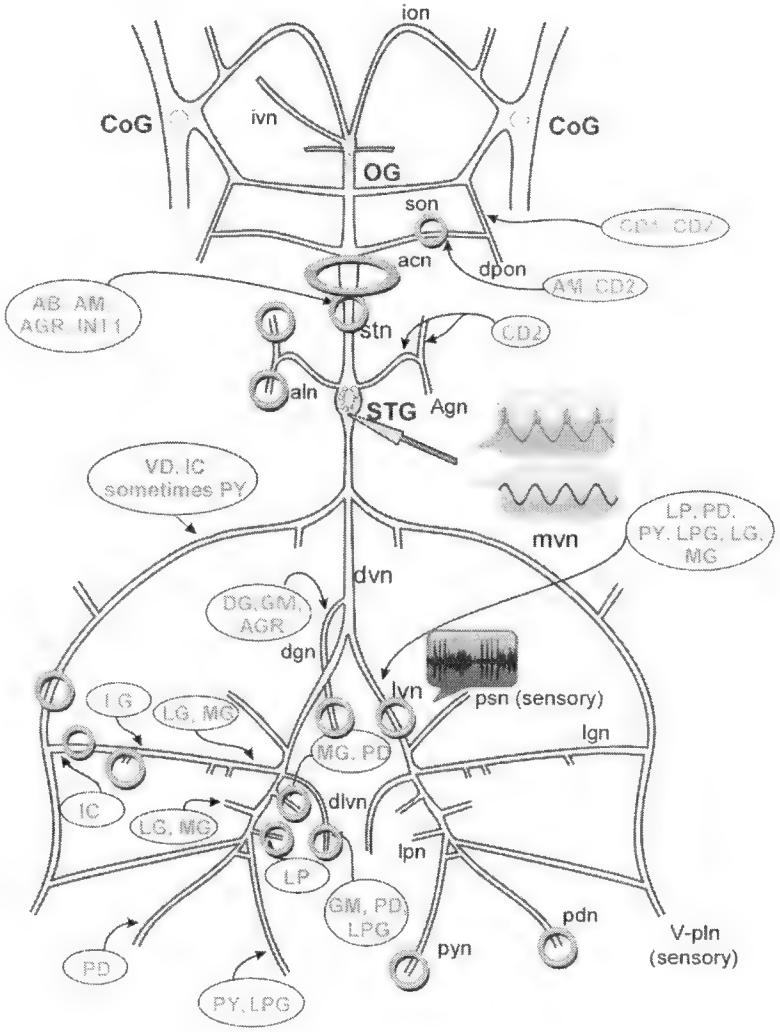
ion
CoG
ivn
CoG
OG
son
CD1, CD2
acn
dpon
AM, CD2
AB, AM,
AGR, INT1
stn
CD2
aln
STG
Agn
VD, IC
sometimes PY
mvn
LP, PD,
PY, LPG, LG,
MG
dvn
DG, GM,
AGR
dgn
lvn
psn (sensory)
lgn
LG
LG, MG
MG, PD
IC
LG, MG
dlvn
lpn
LP
GM, PD,
LPG
pdn
pyn
V-pln
(sensory)
PD
PY, LPG

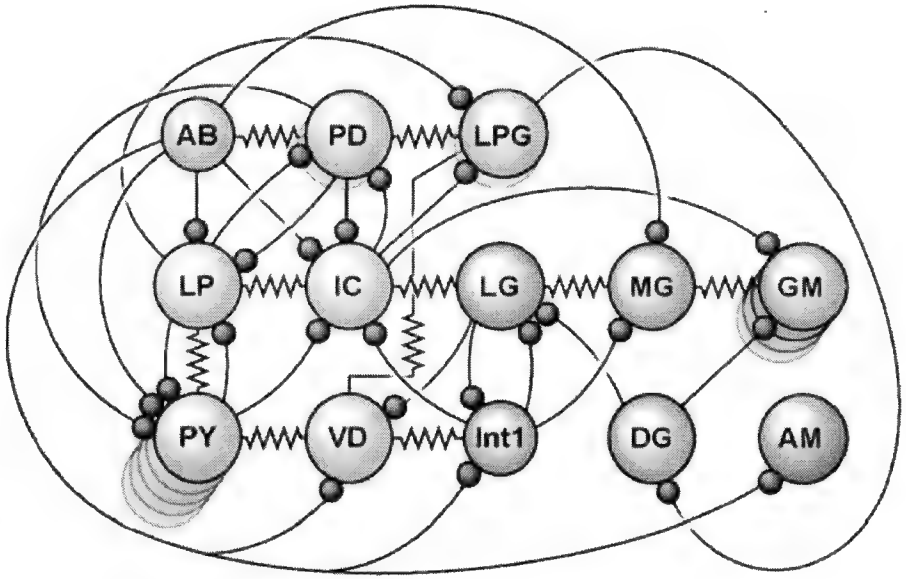
AB
PD
LPG
LP
IC
LG
MG
GM
PY
VD
Int1
DG
AM

$$\left\langle r_i(t) r_j(t) \right\rangle = \sum_{\mu=1}^{N} \lambda_\mu V_i^\mu V_j^\mu$$

$$\left\langle r_i(t) r_j(t) \right\rangle = \sum_{\mu=1}^{N} \lambda_\mu V_i^\mu V_j^\mu$$

$$r_i(t) = \sum_{\mu=1}^{N} \sqrt{\lambda}_\mu c_\mu(t) V_i^\mu$$

$$\left\langle r_i(t) r_j(t) \right\rangle = \sum_{\mu=1}^{N} \lambda_\mu V_i^\mu V_j^\mu$$

$$r_i(t) = \sum_{\mu=1}^{N} \sqrt{\lambda}_\mu c_\mu(t) V_i^\mu$$

then

$$z = \sum_{i=1}^{N} w_i r_i = \sum_{\mu=1}^{\mu_{\max}} a_\mu c_\mu(t)$$

$$\left\langle r_i(t) r_j(t) \right\rangle = \sum_{\mu=1}^{N} \lambda_\mu V_i^\mu V_j^\mu$$

$$r_i(t) = \sum_{\mu=1}^{N} \sqrt{\lambda}_\mu c_\mu(t) V_i^\mu$$

then

$$z = \sum_{i=1}^{N} w_i r_i = \sum_{\mu=1}^{\mu_{\max}} a_\mu c_\mu(t)$$

$$a_\mu = \int dt \; f(t) c_\mu(t)$$

$$z(t) = \int dt' \; f(t')K(t - t')$$

$$z(t) = \int dt' \; f(t')K(t - t')$$

$$K(t - t') = \sum_{\mu=1}^{\mu_{\max}} c_\mu(t)c_\mu(t')$$

$$z(t) = \int dt' \; f(t')K(t - t')$$

$$K(t - t') = \sum_{\mu=1}^{\mu_{\max}} c_\mu(t)c_\mu(t')$$

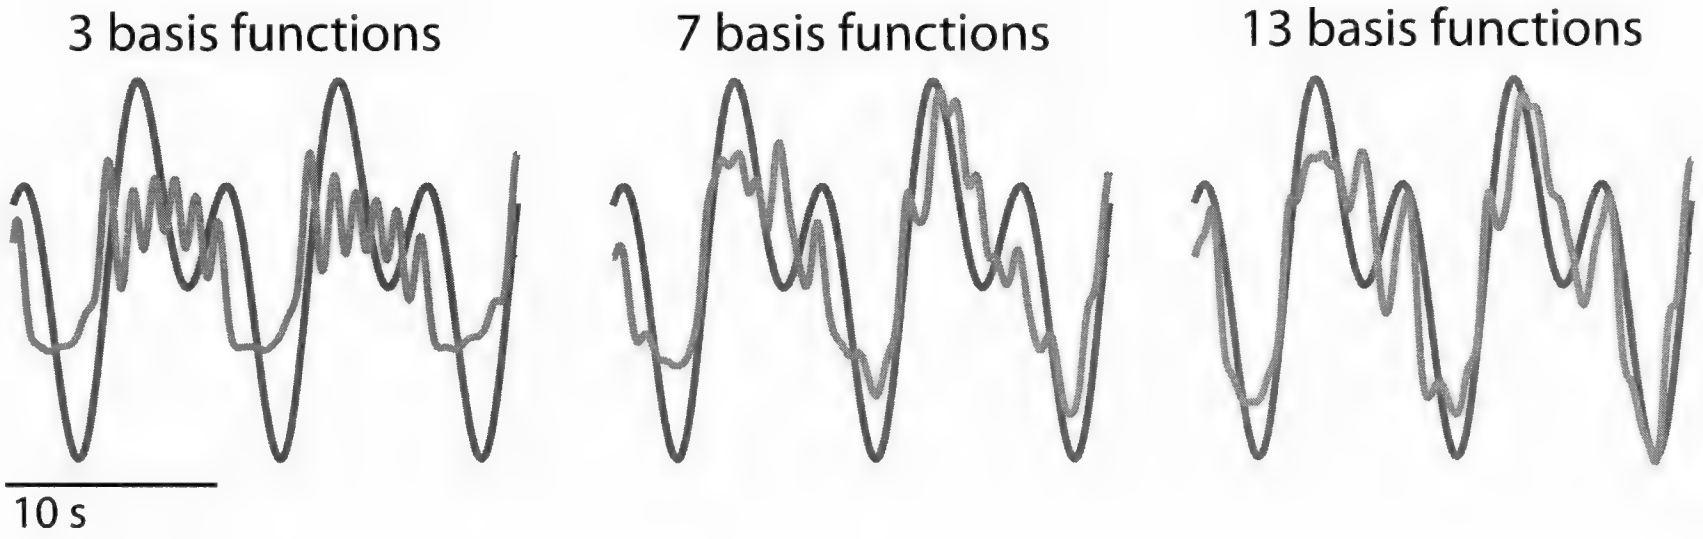
3 basis functions
7 basis functions
13 basis functions
10 s

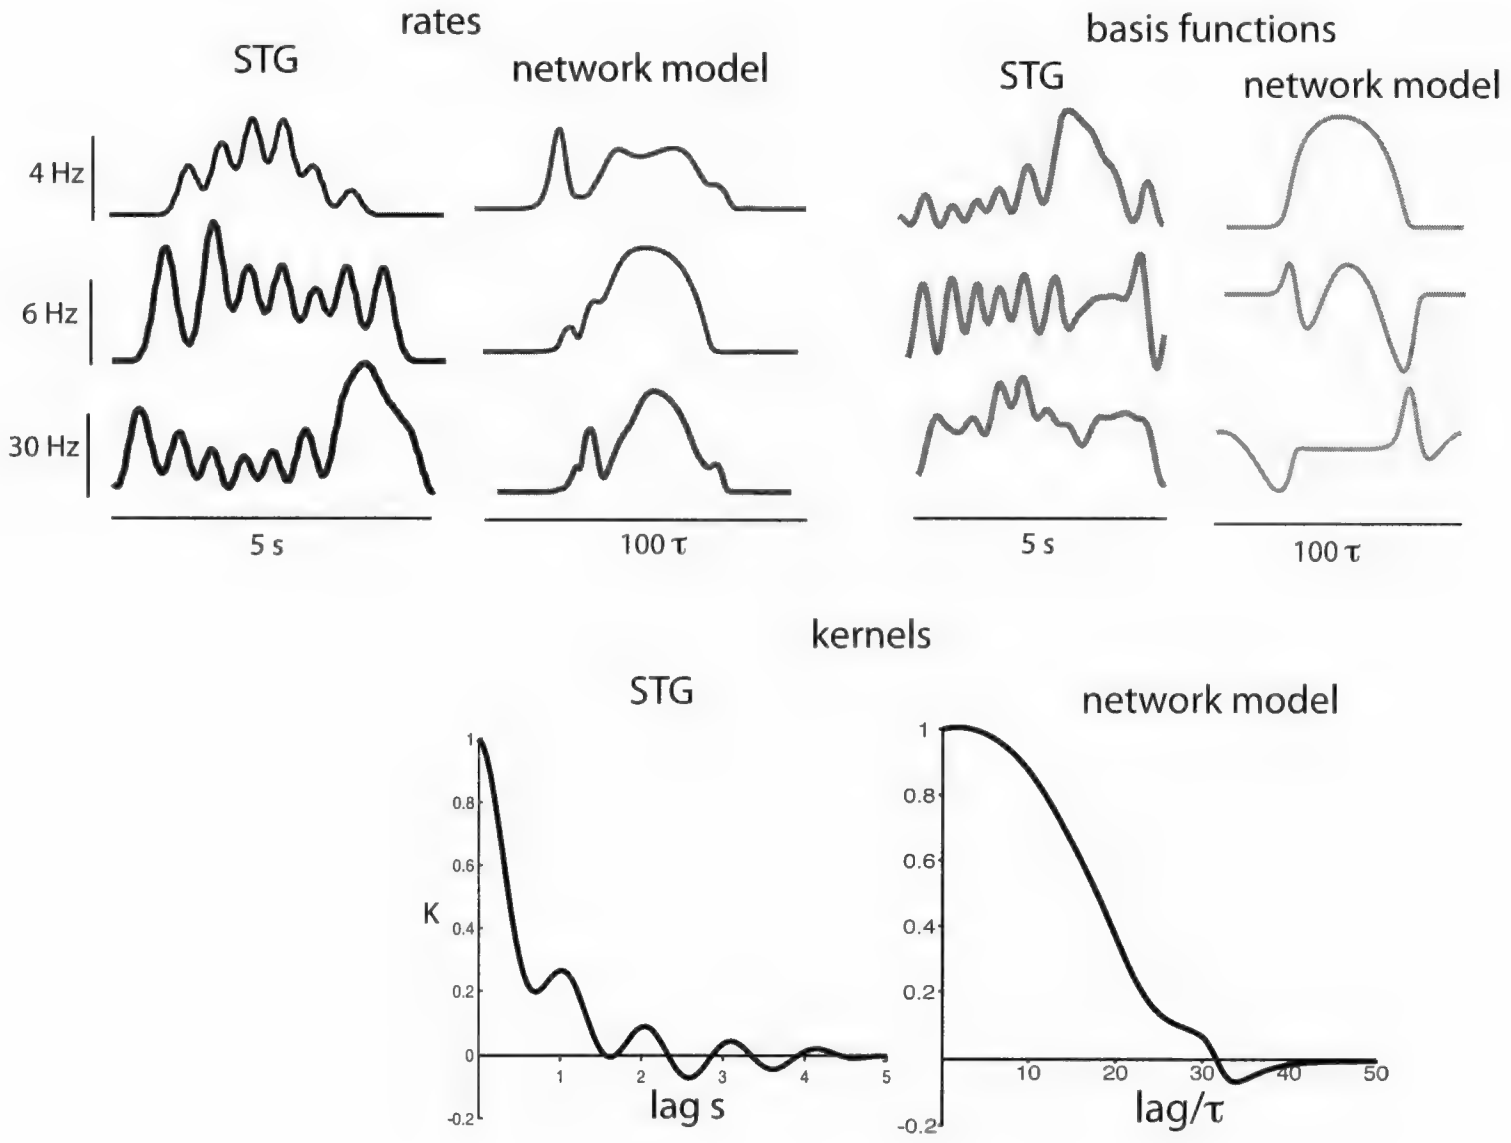
rates
STG
network model
basis functions
STG
network model
4 Hz
6 Hz
30 Hz
5 s
100 τ
5 s
100 τ
kernels
STG
network model
K
lag s
lag/τ

Science
AAAS

Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication

Herbert Jaeger, *et al.*

Science **304**, 78 (2004);
DOI: 10.1126/science.1091277

The following resources related to this article are available online at www.sciencemag.org (this information is current as of October 23, 2009):

Updated information and services, including high-resolution figures, can be found in the online version of this article at:
http://www.sciencemag.org/cgi/content/full/304/5667/78

Supporting Online Material can be found at:
http://www.sciencemag.org/cgi/content/full/304/5667/78/DC1

This article **cites 15 articles**, 4 of which can be accessed for free:
http://www.sciencemag.org/cgi/content/full/304/5667/78#otherarticles

This article has been **cited by** 43 article(s) on the ISI Web of Science.

This article has been **cited by** 6 articles hosted by HighWire Press; see:
http://www.sciencemag.org/cgi/content/full/304/5667/78#otherarticles

Information about obtaining **reprints** of this article or about obtaining **permission to reproduce this article** in whole or in part can be found at:
http://www.sciencemag.org/about/permissions.dtl

Downloaded from www.sciencemag.org on October 23, 2009

Science (print ISSN 0036-8075; online ISSN 1095-9203) is published weekly, except the last week in December, by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. Copyright 2004 by the American Association for the Advancement of Science; all rights reserved. The title *Science* is a registered trademark of AAAS.

Harnessing Nonlinearity: Predicting Chaotic Systems and Saving Energy in Wireless Communication

Herbert Jaeger* and Harald Haas

We present a method for learning nonlinear systems, echo state networks (ESNs). ESNs employ artificial recurrent neural networks in a way that has recently been proposed independently as a learning mechanism in biological brains. The learning method is computationally efficient and easy to use. On a benchmark task of predicting a chaotic time series, accuracy is improved by a factor of 2400 over previous techniques. The potential for engineering applications is illustrated by equalizing a communication channel, where the signal error rate is improved by two orders of magnitude.

Nonlinear dynamical systems abound in the sciences and in engineering. If one wishes to simulate, predict, filter, classify, or control such a system, one needs an executable system model. However, it is often infeasible to obtain analytical models. In such cases, one has to resort to black-box models, which ignore the internal physical mechanisms and instead reproduce only the outwardly observable input-output behavior of the target system.

If the target system is linear, efficient methods for black-box modeling are available. Most technical systems, however, become nonlinear if operated at higher operational points (that is, closer to saturation). Although this might lead to cheaper and more energy-efficient designs, it is not done because the resulting nonlinearities cannot be harnessed. Many biomechanical systems use their full dynamic range (up to saturation) and thereby become lightweight, energy efficient, and thoroughly nonlinear.

Here, we present an approach to learning black-box models of nonlinear systems, echo state networks (ESNs). An ESN is an artificial recurrent neural network (RNN). RNNs are characterized by feedback ("recurrent") loops in their synaptic connection pathways. They can maintain an ongoing activation even in the absence of input and thus exhibit dynamic memory. Biological neural networks are typically recurrent. Like biological neural networks, an artificial RNN can learn to mimic a target system—in principle, with arbitrary accuracy (*1*). Several learning algorithms are known (*2*–*4*) that incrementally adapt the synaptic weights of an RNN in order to tune it toward the target system. These algorithms have not been widely employed in technical applications because of slow convergence and suboptimal solutions (*5*, *6*). The ESN approach differs from these methods in that a large RNN is used (on the order of 50 to 1000 neurons; previous techniques typically use 5 to 30 neurons) and in that only the synaptic connections from the RNN to the output readout neurons are modified by learning; previous techniques tune all synaptic connections (Fig. 1). Because there are no cyclic dependencies between the trained readout connections, training an ESN becomes a simple linear regression task.

We illustrate the ESN approach on a task of chaotic time series prediction (Fig. 2) (*7*). The Mackey-Glass system (MGS) (*8*) is a standard benchmark system for time series prediction studies. It generates a subtly irregular time series (Fig. 2A). The prediction task has two steps: (i) using an initial teacher sequence generated by the original MGS to learn a black-box model M of the generating system, and (ii) using M to predict the value of the sequence some steps ahead.

First, we created a random RNN with 1000 neurons (called the "reservoir") and one output neuron. The output neuron was equipped with random connections that project back into the reservoir (Fig. 2B). A 3000-step teacher sequence $d(1), \ldots, d(3000)$ was generated from the MGS equation and fed into the output neuron. This excited the internal neurons through the output feedback connections. After an initial transient period, they started to exhibit systematic individual variations of the teacher sequence (Fig. 2B).

The fact that the internal neurons display systematic variants of the exciting external signal is constitutional for ESNs: The internal neurons must work as "echo functions" for the driving signal. Not every randomly generated RNN has this property, but it can effectively be built into a reservoir (supporting online text).

It is important that the echo signals be richly varied. This was ensured by a sparse interconnectivity of 1% within the reservoir. This condition lets the reservoir decompose into many loosely coupled subsystems, establishing a richly structured reservoir of excitable dynamics.

After time $n = 3000$, output connection weights w_i ($i = 1, \ldots, 1000$) were computed (dashed arrows in Fig. 2B) from the last 2000 steps $n = 1001, \ldots, 3000$ of the training run such that the training error

$$\mathrm{MSE}_{train} = 1/2000 \sum_{n=1001}^{3000} \left(d(n) - \sum_{i=1}^{1000} w_i x_i(n) \right)^2$$

was minimized [$x_i(n)$, activation of the ith internal neuron at time n]. This is a simple linear regression.

With the new w_i in place, the ESN was disconnected from the teacher after step 3000 and left running freely. A bidirectional dynamical interplay of the network-generated output signal with the internal signals $x_i(n)$ unfolded. The output signal $y(n)$ was created from the internal neuron activation signals $x_i(n)$ through the trained connections w_i, by $y(n) = \sum_{i=1}^{1000} w_i x_i(n)$. Conversely, the internal signals were echoed from that output signal through the fixed output feedback connections (supporting online text).

For testing, an 84-step continuation $d(3001), \ldots, d(3084)$ of the original signal was computed for reference. The network output $y(3084)$ was compared with the correct continuation $d(3084)$. Averaged over 100 independent trials, a normalized root mean square error

$$\mathrm{NRMSE} = \left(\sum_{j=1}^{100} (d_j(3084) - y_j(3084))^2 / 100\sigma^2 \right)^{1/2} \approx 10^{-4.2}$$

was obtained (d_j and y_j teacher and network

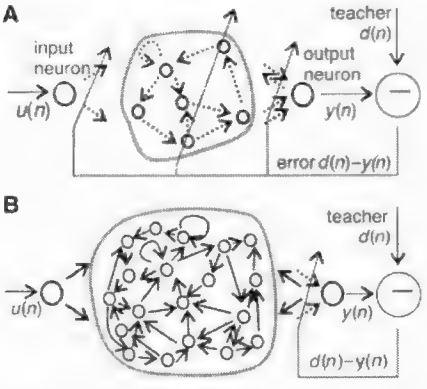

Fig. 1. (**A**) Schema of previous approaches to RNN learning. (**B**) Schema of ESN approach. Solid bold arrows, fixed synaptic connections; dotted arrows, adjustable connections. Both approaches aim at minimizing the error $d(n) - y(n)$, where $y(n)$ is the network output and $d(n)$ is the teacher time series observed from the target system.

International University Bremen, Bremen D-28759, Germany.

*To whom correspondence should be addressed. E-mail: h.jaeger@iu-bremen.de

Downloaded from www.sciencemag.org on October 23, 2009

output in trial j, σ^2 variance of MGS signal), improving the best previous techniques (*9*–*15*), which used training sequences of length 500 to 10,000, by a factor of 700. If the prediction run was continued, deviations typically became visible after about 1300 steps (Fig. 2A). With a refined variant of the learning method (*7*), the improvement factor rises to 2400. Models of similar accuracy were also obtained for other chaotic systems (supporting online text).

The main reason for the jump in modeling accuracy is that ESNs capitalize on a massive short-term memory. We showed analytically (*16*) that under certain conditions an ESN of size N may be able to "remember" a number of previous inputs that is of the same order of magnitude as N. This information is more massive than the information used in other techniques (supporting online text).

We now illustrate the approach in a task of practical relevance, namely, the equalization of a wireless communication channel (*7*). The essentials of equalization are as follows: A sender wants to communicate a symbol sequence $s(n)$. This sequence is first transformed into an analog envelope signal $d(n)$, then modulated on a high-frequency carrier signal and transmitted, then received and demodulated into an analog signal $u(n)$, which is a corrupted version of $d(n)$. Major sources of corruption are noise (thermal or due to interfering signals), multipath propagation, which leads to a superposition of adjacent symbols (intersymbol interference), and nonlinear distortion induced by operating the sender's power amplifier in the high-gain region. To avoid the latter, the actual power amplification is run well below the maximum amplification possible, thereby incurring a substantial loss in energy efficiency, which is clearly undesirable in cell-phone and satellite communications. The corrupted signal $u(n)$ is then passed through an equalizing filter whose output $y(n)$ should restore $u(n)$ as closely as possible to $d(n)$. Finally, the equalized signal $y(n)$ is converted back into a symbol sequence. The quality measure for the entire process is the fraction of incorrect symbols finally obtained (symbol error rate).

To compare the performance of an ESN equalizer with standard techniques, we took a channel model for a nonlinear wireless transmission system from a study (*17*) that compared three customary nonlinear equalization methods: a linear decision feedback equalizer (DFE), which is actually a nonlinear method; a Volterra DFE; and a bilinear DFE. The model equation featured intersymbol interference across 10 consecutive symbols, a second-order and a third-order nonlinear distortion, and additive white Gaussian noise. All methods investigated in that study had 47 adjustable parameters and used sequences of 5000 symbols for training. To make the ESN equalizer comparable with the equalizers studied in (*17*), we took ESNs with a reservoir of 46 neurons (which is small for the ESN approach), which yielded 47 adjustable parameters. (The 47th comes from a direct connection from the input to the output neuron.)

We carried out numerous learning trials (*7*) to obtain ESN equalizers, using an online learning method (a version of the recursive least square algorithm known from linear adaptive filters) to train the output weights on 5000-step training sequences. We chose an online adaptation scheme here because the methods in (*17*) were online adaptive, too, and because wireless communication channels mostly are time-varying, such that an equalizer must adapt to changing system characteristics. The entire learning-testing procedure was repeated for signal-to-noise ratios ranging from 12 to 32 db. Figure 3 compares the average symbol error rates obtained with the results reported in (*17*), showing an improvement of two magnitudes for high signal-to-noise ratios.

For tasks with multichannel input and/or output, the ESN approach can be accommodated simply by adding more input or output neurons (*16*, *18*).

ESNs can be applied to all basic tasks of signal processing and control, including time series prediction, inverse modeling, pattern generation, event detection and classification, modeling distributions of stochastic processes, filtering, and nonlinear control (*16*, *18*, *19*, *20*). Because a single learning run takes only a few seconds (or minutes, for very large data sets and networks), engineers can test out variants at a high turnover rate, a crucial factor for practical usability.

ESNs have been developed from a mathematical and engineering perspective, but exhibit typical features of biological RNNs: a large number of neurons, recurrent pathways, sparse random connectivity, and local modification of synaptic weights. The idea of using randomly connected RNNs to represent and memorize dynamic input in network states has frequently been explored in specific contexts, for instance, in artificial intelligence models of associative memory (*21*), models of prefrontal cortex function in sensory-motor sequencing tasks (*22*), models of birdsong (*23*), models of the cerebellum (*24*), and general computational models of neural oscillators (*25*). Many different learning mechanisms were considered, mostly within the RNN itself. The contribution of the ESN is to elucidate the mathematical properties of large RNNs such that they can be used with a linear, trainable readout mechanism for general black-box modeling. An approach essentially equivalent to ESNs, liquid state networks (*26*, *27*), has been developed independently to model computations in cortical microcircuits. Recent findings in neurophysiology suggest that the basic ESN/liquid state network principle seems not uncommon in biological networks (*28*–*30*) and could eventually be exploited to control prosthetic devices by signals collected from a collective of neurons (*31*).

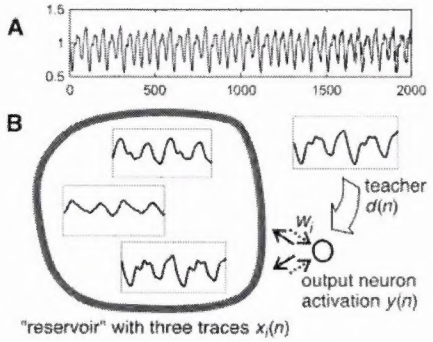

Fig. 2. (**A**) Prediction output of the trained ESN (dotted) overlaid with the correct continuation (solid). (**B**) Learning the MG attractor. Three sample activation traces of internal neurons are shown. They echo the teacher signal $d(n)$. After training, the desired output is recreated from the echo signals through output connections (dotted arrows) whose weights w_i are the result of the training procedure.

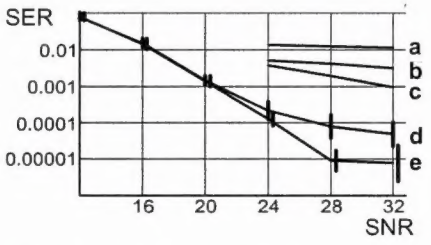

Fig. 3. Results of using an ESN for nonlinear channel equalization. Plot shows signal error rate (SER) versus signal-to-noise ratio (SNR). (**a**) Linear DFE. (**b**) Volterra DFE. (**c**) Bilinear DFE. [(a) to (c) taken from (*20*)]. (**d**) Blue line represents average ESN performance with randomly generated reservoirs. Error bars, variation across networks. (**e**) Green line indicates performance of best network chosen from the networks averaged in (d). Error bars, variation across learning trials.

References and Notes

1. K.-I. Funahashi, Y. Nakamura, *Neural Netw.* **6**, 801 (1993).
2. D. Zipser, R. J. Williams, *Neural Comput.* **1**, 270 (1989).
3. P. J. Werbos, *Proc. IEEE* **78**, 1550 (1990).
4. L. A. Feldkamp, D. V. Prokhorov, C. F. Eagen, F. Yuan, in *Nonlinear Modeling: Advanced Black-Box Techniques*, J. A. K. Suykens, J. Vandewalle, Eds. (Kluwer, Dordrecht, Netherlands, 1998), pp. 29–54.
5. K. Doya, in *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, Ed. (MIT Press, Cambridge, MA, 1995), pp. 796–800.
6. H. Jaeger, "Tutorial on training recurrent neural networks" (GMD-Report 159, German National Research Institute for Computer Science, 2002); ftp://borneo.gmd.de/pub/indy/publications_herbert/CompleteTutorialTechrep.pdf.

Downloaded from www.sciencemag.org on October 23, 2009

7. Materials and methods are available as supporting material on *Science* Online.
8. M. C. Mackey, L. Glass, *Science* **197**, 287 (1977).
9. J. Vesanto, in *Proc. WSOM '97* (1997); www.cis.hut.fi/projects/monitor/publications/papers/wsom97.ps.
10. L. Chudy, I. Farkas, *Neural Network World* **8**, 481 (1998).
11. H. Bersini, M. Birattari, G. Bontempi, in *Proc. IEEE World Congr. on Computational Intelligence* (IJCNN '98) (1997), pp. 2102–2106; ftp://iridia.ulb.ac.be/pub/lazy/papers/IridiaTr1997-13_2.ps.gz.
12. T. M. Martinetz, S. G. Berkovich, K. J. Schulten, *IEEE Trans. Neural Netw.* **4**, 558 (1993).
13. X. Yao, Y. Liu, *IEEE Trans. Neural Netw.* **8**, 694 (1997).
14. F. Gers, D. Eck, J. F. Schmidhuber, "Applying LSTM to time series predictable through time-window approaches" (IDSIA-IDSIA-22-00, 2000); www.idsia.ch/~felix/Publications.html.
15. J. McNames, J. A. K. Suykens, J. Vandewalle, *Int. J. Bifurcat. Chaos* **9**, 1485 (1999).
16. H. Jaeger, "Short term memory in echo state networks" (GMD-Report 152, German National Research Institute for Computer Science, 2002); ftp://borneo.gmd.de/pub/indy/publications_herbert/STMEchoStatesTechRep.pdf.
17. V. J. Mathews, J. Lee, in *Advanced Signal Processing: Algorithms, Architectures, and Implementations V (Proc. SPIE Vol. 2296)*, (SPIE, San Diego, CA, 1994), pp. 317–327.
18. J. Hertzberg, H. Jaeger, F. Schönherr, in *Proc. 15th Europ. Conf. on Art. Int. (ECAI 02)*, F. van Harmelen, Ed. (IOS Press, Amsterdam, 2002), pp. 708–712; www.ais.fhg.de/~schoenhe/papers/ECAI02.pdf.
19. H. Jaeger, "The echo state approach to analysing and training recurrent neural networks" (GMD-Report 148, German National Research Institute for Computer Science, 2001); ftp://borneo.gmd.de/pub/indy/publications_herbert/EchoStatesTechRep.pdf.
20. H. Jaeger, in *Advances in Neural Information Processing Systems* 15, S. Becker, S. Thrun, K. Obermayer, Eds. (MIT Press, Cambridge, MA, 2003) pp. 593–600.
21. G. E. Hinton, in *Parallel Models of Associative Memory*, G. E. Hinton, J. A. Anderson, Eds. (Erlbaum, Hillsdale, NJ, 1981), pp. 161–187.
22. D. G. Beiser, J. C. Houk, *J. Neurophysiol.* **79**, 3168 (1998).
23. S. Dehaene, J.-P. Changeux, J.-P. Nadal, *Proc. Natl. Acad. Sci. U.S.A.* **84**, 2727 (1987).
24. M. Kawato, in *The Handbook of Brain Theory and Neural Networks*, M. Arbib, Ed. (MIT Press, Cambridge, MA, 1995), pp. 172–178.
25. K. Doya, S. Yoshizawa, *Neural Netw.* **2**, 375 (1989).
26. W. Maass, T. Natschläger, H. Markram, *Neural Comput.* **14**, 2531 (2002).
27. W. Maass, T. Natschläger, H. Markram, in *Computational Neuroscience: A Comprehensive Approach*, J. Feng, Ed. (Chapman & Hall/CRC, 2003), pp. 575–605.
28. G. B. Stanley, F. F. Li, Y. Dan, *J. Neurosci.* **19**, 8036 (1999).
29. G. B. Stanley, *Neurocomputing* **38–40**, 1703 (2001).
30. W. M. Kistler, Ch. I. de Zeeuw, *Neural Comput.* **14**, 2597 (2002).
31. S. Mussa-Ivaldi, *Nature* **408**, 361 (2000).

32. The first author thanks T. Christaller for unfaltering support and W. Maass for friendly cooperation. International patents are claimed by Fraunhofer AIS (PCT/EP01/11490).

Supporting Online Material

www.sciencemag.org/cgi/content/full/304/5667/78/DC1
Materials and Methods
SOM Text
Figs. S1 to S4
References

8 September 2003; accepted 26 February 2004

Ultrafast Electron Crystallography of Interfacial Water

Chong-Yu Ruan, Vladimir A. Lobastov, Franco Vigliotti, Songye Chen, Ahmed H. Zewail*

We report direct determination of the structures and dynamics of interfacial water on a hydrophilic surface with atomic-scale resolution using ultrafast electron crystallography. On the nanometer scale, we observed the coexistence of ordered surface water and crystallite-like ice structures, evident in the superposition of Bragg spots and Debye-Scherrer rings. The structures were determined to be dominantly cubic, but each undergoes different dynamics after the ultrafast substrate temperature jump. From changes in local bond distances (OH···O and O···O) with time, we elucidated the structural changes in the far-from-equilibrium regime at short times and near-equilibration at long times.

The nature of interfacial molecular assemblies of nanometer scale is of fundamental importance to chemical and biological phenomena (*1*–*4*). For water, the directional molecular features of hydrogen bonding (*5*, *6*) and the different structures possible, from amorphous (*7*) to crystalline (*8*), make the interfacial (*9*) collective assembly on the mesoscopic (*10*) scale much less understood. Structurally, the nature of water on a substrate is determined by forces of orientation at the interface and by the net charge density, which establishes the hydrophilic or hydrophobic character of the substrate. However, the transformation from ordered to disordered structure and their coexistence critically depends on the time scales for the movements of atoms locally and at long range. Therefore, it is essential to elucidate the nature of these structures and the time scales for their equilibration.

Here, we report direct determination of the structures of interfacial water with atomic-scale resolution, using diffraction and the dynamics following ultrafast infrared (IR) laser-initiated temperature jump. Interfacial water is formed on a hydrophilic surface (silicon, chlorine-terminated) under controlled ultrahigh vacuum (UHV) conditions (Fig. 1). With these atomic-scale spatial, temporal, and energy resolutions, the evolution of nonequilibrium structures was monitored, their ordered or disordered nature was established, and the time scale for the breakage of long-range bonding and formation of new structures was determined. We identified the structured and ordered interfacial water from the Bragg diffraction and the layered crystallite structure from the Debye-Scherrer rings. The temporal evolution of interfacial water and layered ice after the temperature jump was studied with submonolayer sensitivity. We compared these results with those obtained on hydrophobic surfaces, such as hydrogen-terminated silicon or silver substrate.

Spectroscopic techniques, such as internal reflection (*11*) and nonlinear [second-harmonic generation (*12*) and sum-frequency generation

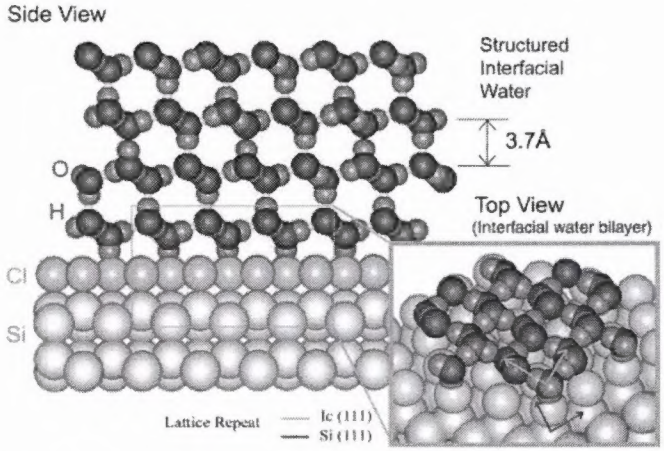

Fig. 1. Structured water at the hydrophilic interface. The chlorine termination on a Si(111) substrate forms a hydrophilic layer that orients the water bilayer. The closest packing distance (4.43 Å) between oxygen atoms in the bottom layer of water is similar to the distance (4.50 Å) between the on-top and interstitial sites of the chlorine layer, resulting in specific bilayer orientations (±30°) with respect to the silicon substrate. This ordered stacking persists for three to four bilayers (~1 nm) before disorientation takes place and results in crystallite islands, forming the layered structure. The size of atoms is not to scale for the van der Waals radii.

Laboratory for Molecular Sciences, Arthur Amos Noyes Laboratory of Chemical Physics, California Institute of Technology, Pasadena, CA 91125, USA.

*To whom correspondence should be addressed. E-mail: zewail@caltech.edu

Downloaded from www.sciencemag.org on October 23, 2009